Liquid Cooling Guidelines for Datacom Equipment Centers

ASHRAE Datacom Series

American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

Liquid Cooling Guidelines for Datacom Equipment Centers

This publication was prepared in cooperation with TC 9.9, Mission Critical Facilities, Technology Spaces, and Electronic Equipment.

Any updates/errata to this publication will be posted on the ASHRAE Web site at www.ashrae.org/publicationupdates.

Liquid Cooling Guidelines for Datacom Equipment Centers

American Society of Heating, Refrigerating
and Air-Conditioning Engineers, Inc.

ISBN-10: 1-933742-05-4
ISBN-13: 978-1-933742-05-2

1791 Tullie Circle, NE
Atlanta, GA 30329
www.ashrae.org

Printed in the United States of America

Printed on 30% post-consumer waste using soy-based inks.

Library of Congress Cataloging-in-Publication Data

Liquid cooling guidelines for datacom equipment centers.
p. cm.
Includes bibliographical references and index.
ISBN 978-1-933742-05-2 (pbk.)
1. Data processing service centers--Cooling. 2. Hydronics. 3. Electronic apparatus and appliances--Cooling. 4. Electronic digital computers--Cooling. I. American Society of Heating, Refrigerating and Air-Conditioning Engineers.

TH7688.C64L57 2006
697.9'316--dc22
2006026115

ASHRAE Staff

Contents

Acknowledgments

Representatives from the following companies participated in producing this publication:

APC
Aavid
Cray Inc.
Dell Computers
DLB Associates Consulting Engineers
EYPMCF
Hewlett Packard
IBM
Intel Corporation
Lawrence Berkeley National Labs
Liebert Corporation
Lytron
Mallory & Evans, Inc.
NCR
Department of Defense
Panduit
Rittal
Sanmina
SGI
Spraycool
Syska & Hennessy Group, Inc.
Sun Microsystems
Trane

ASHRAE TC 9.9 wishes to particularly thank the following people:

- **Don Beaty, John Bean, Christian Belady, Tahir Cader, David Copeland, Rhonda Johnson, Tim McCann, David Moss, Shlomo Novotny, Greg Pautsch, Terry Rodgers, Jeff Rutt, Tony Sharp, David Wang,** and **Kathryn Whitenack** for their participation and continual improvements in the final document.
- **Don Beaty** of DLB Associates Consulting Engineers, **Jeff Rutt** of Department of Defense, and **Kathryn Whitenack** of Lytron for providing the coordination and leadership on developing the individual chapters.

- **Dr. Roger Schmidt** for his invaluable participation in leading the overall development of the book.

In addition, ASHRAE TC 9.9 wishes to thank the following people for their comments and feedback: William Angle, Cullen Bash, Neil Chauhan, Brian Donabedian, Jack Glass, Chris Malone, Vance Murakami, Larry Rushing, Prajbit Singh, William Tschudi, and Randy Zoodsma.

1

Introduction

From a holistic point of view, the data center, with its installed information technology (IT) equipment, introduces several levels of cooling. First there is the utility plant that provides the overall cooling to the data center. This utility plant may be a stand-alone building dedicated solely to the data center or, as is more typically the case, part of a larger building that provides cooling to other locations within the building where the data center is housed. The utility plant that provides the cooling to the building/data center generally employs large chillers that are used to chill the building's water loop. The chilled water can be used to provide cooling to modular computer room air-conditioning (CRAC) units situated inside a data center or to large air-handling units installed externally to the data center (even to the building). Air conditioning for the data center can also be provided by refrigeration units, e.g., DX units or expansion type devices, whereby the condenser unit is generally placed outside the building envelope. In this case, no water is delivered to the data center floor.

Data center IT equipment today is generally air-cooled by the means described in the preceding paragraph. With rack heat loads steadily climbing, the ability for many data centers to deliver either adequate airflow rates or sufficient chilled air is now being stretched to the limit. This situation is creating a need for implementing liquid cooling solutions. The overall goals of the liquid implementations are to transfer as much waste heat to the facility water as possible and, in some of the implementations, to reduce the overall volume of airflow needed by the racks. In addition, implementation of liquid cooling may be required to achieve higher performance of the datacom equipment through lower temperatures achieved with the cooling of microprocessors.

This book, which was generated by ASHRAE Technical Committee 9.9, provides equipment manufacturers and facility operations personnel with a common set of guidelines for liquid cooling. This publication is not inclusive of all types of liquid cooling sources (e.g., absorption chillers) but is representative of generally accepted liquid cooling systems. It covers an overview of liquid cooling, various liquid cooling configurations, and guidelines for liquid cooling infrastructure requirements. Specifically, this book provides guidelines for:

- Chilled-water operational requirements (Section 5.1.1.1)
- Chilled-water flow rate recommendations and fouling factors (Section 5.1.1.2)
- Chilled-water velocity considerations (Section 5.1.1.3)
- Chilled-water liquid quality/composition (Section 5.1.1.4)
- Chilled-water wetted material requirements (Section 5.1.1.5)
- Refrigerant piping (Section 5.2.2)
- Rack-water operational requirements (Section 6.1.1)
- Rack-water flow rates (Section 6.1.2)
- Rack-water velocity considerations (Section 6.1.3)
- Rack-water quality/composition (Section 6.1.4)
- Rack-water wetted materials requirements (Section 6.1.5)
- Rack-nonwater operational requirements (Section 6.2.1)

It also reviews considerations for:

- Chilled-water piping
- Electrical power sources and connections
- Monitoring
- Reliability and availability
- Commissioning

The following chapters are arranged to describe the various liquid cooling systems that can exist within a data center and, more importantly, how they can be connected to transfer heat from one liquid system to another. Chapters 2 and 3 provide the reader with an overview of the chilled-water and condenser-water systems (see Figure 1.1), followed by an overview of datacom equipment cooling options (technology cooling system and datacom equipment cooling system shown in Figure 1.1) in Chapter 4. Chapter 5 bridges the liquid cooling systems by providing guidelines on the interface requirements between the chilled-water system and the technology cooling system. Chapter 6 outlines the requirements of those liquid-cooled systems that attach to a datacom electronics rack and are implemented to aid in data center thermal management by focusing on the technology cooling system. Chapter 7 provides an expanded reference and bibliography list to enable the reader to find additional related materials. Chapter 8 provides a useful glossary of common terms used throughout this book. An appendix is also included that provides survey results of customer water quality.

1.1 DEFINITIONS

For the purposes of this book, the following definitions will be used:

- *Liquid cooling* is defined as the case where liquid must be supplied to an entity for operation.

- *Liquid-cooled rack* defines the case where liquid must be circulated to and from the rack or cabinet for operation.
- *Liquid-cooled datacom equipment* defines the case where liquid must be circulated to and from the datacom equipment for operation.
- *Liquid-cooled electronics* defines the cases where liquid is supplied directly to the electronics for cooling with no other form of heat transfer.

In each of the definitions above, when two-phase liquid cooling is employed, liquid is circulated to the rack, equipment, or electronics, and gas and/or a mixture of gas and liquid circulates from the rack, equipment, or electronics. These liquid cooling definitions are slightly broader than the one used in ASHRAE's *Datacom Equipment Power Trends and Cooling Applications* but are more relevant to liquid cooling in the datacom industry today.

It is important to keep in mind that the definitions above do not limit the cooling fluid to water. A variety of liquids could be considered for application, including liquids that could be in a vapor phase in part of the cooling loop.

- *Air cooling* defines the case where only air must be supplied to an entity for operation.
- *Air-cooled rack* defines the case where only air must be provided to the rack or cabinet for operation.
- *Air-cooled datacom equipment* defines the case where only air is provided to the datacom equipment for operation.
- *Air-cooled electronics* defines the cases where air is provided directly to the electronics for cooling with no other form of heat transfer.

When liquids are employed within separate cooling loops that do not communicate thermally, the system is considered to be air cooling. The most obvious illustration covers the chilled-water CRACs that are usually deployed at the periphery of many of today's data centers. At the other end of the scale, the use of heat pipes or pumped loops inside a computer, wherein the liquid remains inside a closed loop within the server, also qualifies as air-cooled electronics, provided the heat is removed from the internal closed loop via airflow through the electronic equipment chassis.

There are many different implementations of liquid cooling to choose from. Below are several scenarios (see Chapter 4 for a more complete overview of liquid cooling implementations).

- One option uses an air-cooled refrigeration system mounted within the datacom equipment to deliver chilled refrigerant to liquid-cooled cold plates mounted to the processors. For this implementation, the heated air from the liquid-to-air heat exchanger (i.e., condenser) is exhausted directly to the data center environment. From a data center perspective, the rack and electronics are considered to be air-cooled since no liquid lines cross the rack envelope.

- A different implementation may use a liquid-to-air heat exchanger mounted above, below, or on the side or rear of the rack. In this case, the heat exchanger removes a substantial portion of the rack's waste heat from the air that is eventually exhausted to the data center. This implementation does not reduce the volumetric airflow rate needed by the electronics, but it does reduce the temperature of the air that is exhausted back into the data center. This example describes a liquid-cooled rack since liquid lines cross the rack envelope. This system is shown in Figure 4.8.
- Yet another implementation uses liquid-cooled cold plates that employ water, dielectrics, or other types of coolants that are chilled by a liquid-to-liquid heat exchanger that rejects the waste heat to the facility water. The waste heat rejection to the facility water can occur via one or more additional liquid loops that eventually terminate at an external cooling tower or chiller plant. This implementation of liquid cooling reduces the amount of waste heat rejected to the facility ambient and also reduces the volumetric airflow rate required by the rack's electronics. From the data center perspective, this implementation describes liquid-cooled racks and electronics since liquid lines cross the rack envelope and also cross over into the servers themselves. This system is shown in Figure 4.10.

1.2 LIQUID COOLING SYSTEMS

The definitions above apply to air and liquid cooling in a data center. They do not define the various liquid cooling loops that may exist within a data center. Figure 1.1 shows a typical liquid-cooled facility with multiple liquid cooling loops. Each loop is described below with some of the variations considered. The terminology used below will be applied throughout the rest of the book.

datacom equipment cooling system (DECS): This system does not extend beyond the IT rack. It is a loop within the rack that is intended to perform heat transfer from the heat-producing components (CPU, memory, power supplies, etc.) to a fluid-cooled heat exchanger also contained within the IT rack. Some configurations may eliminate this loop and have the fluid from the coolant distribution unit (CDU) flow directly to the load. This loop may function in single-phase or two-phase heat transfer modes facilitated by heat pipes, thermosyphon, pumped fluids, and/or vapor-compression cycles. Fluids typically used in the datacom equipment include water, ethylene glycol or propylene glycol and water mixture, refrigerants, or dielectrics. At a minimum the datacom equipment cooling system would include a heat collection heat exchanger as well as a heat-of-rejection heat exchanger and may be further enhanced with active components such as compressor/pump, control valves, electronic controls, etc.

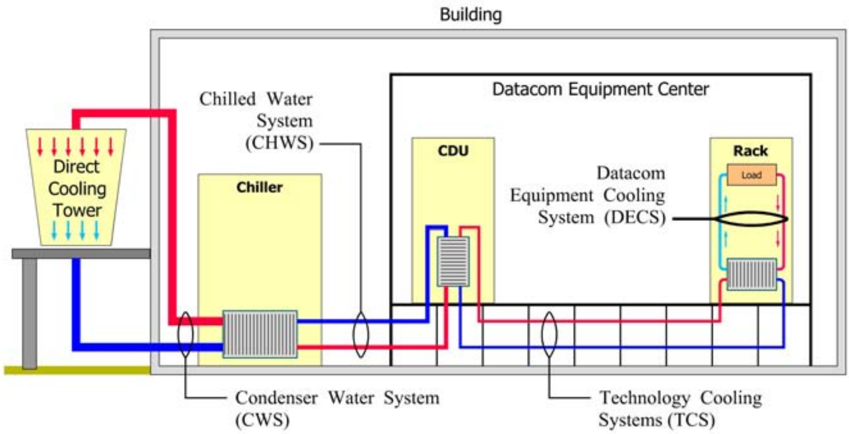

Figure 1.1 Liquid cooling systems/loops within a data center.

technology cooling system (TCS): This system would not typically extend beyond the boundaries of the IT space. The exception is a configuration in which the CDU is located outside the data center. It serves as a dedicated loop intended to perform heat transfer from the datacom equipment cooling system into the chilled-water system. This loop is highly recommended, as it is needed to address specific fluid quality issues regarding temperature, purity, and pressure as required by the heat exchangers within the datacom equipment cooling systems. Fluids typically used in the technology cooling loop include water, ethylene glycol or propylene glycol and water mixture, refrigerants, or dielectrics. This loop may also function by single-phase or two-phase heat transfer modes and may facilitate transfer by heat pipes, thermosyphon, pumped fluids, and/or vapor-compression cycles. At a minimum the technology cooling system would include a heat collection heat exchanger (likely integral component of the datacom equipment cooling system), a heat rejection heat exchanger, as well as interconnecting piping. This system may be further enhanced with such active components as compressors/pumps, control valves, electronic controls, filters, hydronic accessories, etc.

chilled-water system (CHWS): This system is typically at the facility level and may include a dedicated system for the IT space(s). It primarily consists of the system between the data center chiller(s) and the CDU. The chilled-water system would include the chiller plant, pumps, hydronic accessories, and necessary distribution piping at the facility level. The chiller plant would typically employ a vapor-compression cycle to cool the chilled-water supply temperature (43°F–48°F/6°C–

9°C) substantially below indoor ambient temperature (typically 75°F/24°C and up to and beyond 95°F/35°C). The chiller system may offer some level of redundancy for critical components such as chillers, cooling towers, and pumps.

DX equipment can also be used in the chilled-water system. DX equipment provides direct heat dissipation to the atmosphere and is therefore the last loop for that design method. Limitations include distance for the split systems and cost of operation. Generally, in most areas systems become economically breakeven at 400 tons of refrigeration. Larger systems favor non-DX designs unless other circumstances warrant more extensive DX deployment. Smaller thermal ride-through devices can be introduced for individual or special cases within this loop design.

condenser-water system (CWS): This system consists of the liquid loop between the cooling towers and the data center chiller(s). It is also typically at the facility level and may or may not include a dedicated system for the IT space(s). Condenser-water loops typically fall into one of two fundamental categories: wet-bulb-based or dry-bulb-based system. The wet-bulb-based loops function on an evaporative process, taking advantage of lower wet-bulb temperatures, thereby providing cooler condenser-water temperatures. The dry-bulb-based loops function based upon the difference of condenser-water loop temperature versus ambient dry-bulb temperature. To allow heat transfer with the dry-bulb-based system, the condenser-water loop must be at some temperature substantially above the ambient dry-bulb temperature to allow adequate heat transfer from the condenser-water into the outdoor ambient air. These loops would typically include: outdoor heat rejection device (cooling tower or dry fluid cooler), pumps, expansion tanks, hydronic accessories, and distribution piping.

Any of these systems can be eliminated. For instance, in an installation where the CDU fluid flows directly to the load, the datacom equipment cooling system is eliminated or, where fluid from the chilled-water system flows to a heat exchanger in the rack (this is strongly discouraged), the technology cooling system is eliminated or the technology cooling system and datacom equipment cooling system are eliminated (again, this is strongly discouraged).

General ranges for air-cooled equipment and evaporative water cooling as well as the refrigerant types and compressor efficiencies result in a unique engineering design for each data center. For cooling component selection, the design engineer must consider economical operation, first cost, expandability, reliability, redundancy, and fault tolerance. To complicate the design further, all of the systems are interrelated. For example, undersize one loop and that will limit the successive removal of heat to the next level. When designing a data center, the interdependence must be evaluated. ASHRAE's Datacom Book Series (specifically, *Design Considerations for Datacom Equipment Centers*, *Datacom Equipment Power Trends and Cooling Applications*, and *Thermal Guidelines for Data Processing Environments*) provides more details about the various systems and their design requirements.

2

Facility Cooling Systems

2.1 INTRODUCTION

Facility cooling systems can vary in configuration and equipment. Time and cost pressures can cause a disproportionate focus on simply designing cooling systems that deliver the capacity needed and ignoring other critical considerations. However, the focus should be much broader and balance the various considerations and trade-offs, such as flexibility; scalability; ease of installation, commissioning, and operation; ease of maintenance and troubleshooting; and availability and reliability. Any one or combination of the five considerations just mentioned can significantly change how the facility cooling systems are designed, what equipment is used, and the overall system architecture.

2.1.1 Flexibility

Data center cooling systems should be designed with features that will minimize or eliminate system outages associated with new equipment installation. These features should be added to both the central plant cooling systems and building chilled-water piping architecture. Some of these features include valved and capped piping connections for future equipment, such as water-cooled racks, central station air handlers, CRACs, and central plant equipment. The central plant should be configured to add additional chillers, pumps, and cooling towers as the load increases. A properly managed load and growth plan or strategy should be developed and employed to incorporate future computer and cooling systems.

Overall flexibility is often limited by the pipe sizes used in the central plant and distribution system. After a data center is online, changing pipe size to increase capacity is typically prohibitive from both outage risk and implementation cost perspectives.

2.1.2 Scalability

The building cooling systems should be designed to accommodate future load growth of the computer equipment. Unless adequately planned growth and expansion is included in the data center, it will be obsolete in a very short time. Computer technology changes every two to five years, and the cooling system will need to be expanded to accommodate this load growth. Therefore, building piping architecture (CWS and CHWS) should be designed to support a future building cooling load density. Although first cost often dictates pump selection and pipe sizing, pumping energy, flexibility, and chilled-water storage will need to be considered to determine a total cost of ownership.

The central plant should have enough space for future chillers, pumps, and cooling towers. The central plant chilled- and condenser-water system pipe headers should be sized to operate efficiently from day one through the ramp up and for the future projected load.

Sizing piping for the full buildout or future growth will save energy and allow smaller active components during the early life of the building. If the budget does not allow for any additional building costs for future growth, the owner needs to make sure that real estate is available next to the existing central plant.

2.1.3 Ease of Installation, Commissioning, and Operation

Cooling service equipment should be designed such that it can be installed in easy, visible, and readily accessible locations.

Commissioning is an effective strategy to verify that the cooling systems are operating as intended in the original building design and should be considered for every project. In *Design Considerations for Datacom Equipment Centers* (ASHRAE 2005a), Chapter 12 provides helpful information regarding cooling services for equipment cooling systems. This chapter details the five steps of formal commissioning activities, starting with the facility's intent and performance requirements (as determined by the project team), and following with the Owner's Program document, the Basis-of-Design document, and the project Commissioning Plan. These activities include factory acceptance tests, field component verification, system construction verification, site acceptance testing, and integrated systems testing.

Commissioning at full load (full flow) to prove hydraulic capacity should be a requirement. Loop isolation segments should be commissioned to prove circulation capacity around each segment with looped supply and return to provide concurrent maintainability without loss of cooling service.

Data centers should be designed with a central control or command center to oversee building operations. The control center should house all of the building operations systems, such as security monitoring, energy management and control systems (EMCS), system control and data acquisition (SCADA) systems, building automation systems (BAS), and fire alarms. This control can be co-located with the computer system control room. It should be staffed for 24-hour operation. All emergency procedures, protocols, and a personnel list should be included in the control center and updated as changes occur. Power and cooling loads on facility equipment, such as UPS, chillers, and electrical feeders, should be monitored to determine load growth and available capacity. Configuration management groups or boards, consisting of IT and facilities departments, should be established to control and manage data center infrastructure.

2.1.4 Ease of Maintenance and Troubleshooting

The ease of maintenance and the ability to troubleshoot problems quickly and accurately are essential elements of a high-availability datacom facility. The first element of this planning should be to maintain adequate working clearance around cooling equipment. Manufacturers' recommendations for working clearances should be used as a minimum for serviceability areas. Designers need to provide access to allow maintenance and operation of valves, control devices and sensors, and large equipment. Lifts, hoists, and cranes may be mounted within the central plant to help facilitate removal of heavy equipment and components. These devices should be incorporated into the existing structural members of the plant room. For example, a hoist and rail system could be placed above the chillers in the plant room to facilitate rapid removal of end plates and/or compressors. Also, space should be provided to facilitate demolition and removal of entire cooling system components, such as chillers. If space is at a premium, tube pull areas for each chiller can be shared since usually one chiller is being worked on at a time. Chilled and condenser water pipes should be routed to avoid conflict with removal of cooling system equipment. Mechanical equipment, such as pumps and chillers, should be arranged in such a manner as to facilitate complete replacement. Isolation valves must also be located to allow for replacement without interrupting service, which makes layout and assembly of the entire piping system important.

Building owners should consider a computerized maintenance management system (CMMS) to help manage equipment maintenance. These systems can record maintenance history and automatically dispatch work orders for future maintenance. Manufacturers' specific maintenance requirements and frequencies can be input or downloaded into the CMMS. It is much easier and desirable to coordinate equipment outages and maintenance than to deal with an unscheduled equipment failure due to lack of adequate maintenance.

Energy management and control systems (EMCS) or building automation systems (BAS) sensors and device outputs can be trended over time and used for system diagnostics and troubleshooting. EMCS data can also be used to monitor and characterize system performance over time. For example, chiller amp readings and/or chilled-water temperature differentials and flows can be used to calculate and monitor chiller efficiency and load growth. Control systems should have a fail-safe condition that allows mechanical flow and ON operation. Typical general building management systems shut down on the loss of control. Data center systems should turn on to keep systems online.

Load growth over time can be compared to chilled-water capacity, and this information can be used to project time frames for plant expansion or increase in individual component capacity, i.e., replace an 800 ton chiller with a 1200 ton chiller. In addition, flowmeters and pressure sensors can be installed in the building cooling distribution or secondary loop and used to monitor chilled-water flows and capacities in the piping architecture. This information can be used to determine the best location in which to install the newest water-cooled computer equipment or used to calibrate a network model of the chilled-water systems. Finally, analog thermometers, pressure gauges, and flow-measuring instrumentation (orifice plates, balancing valves, etc.) should be installed in chilled-water piping and used to gain additional information on system performance. Sensors need placement in both the primary loop and the auxiliary loop to allow control if either part of the loop is being serviced. Manual operation may also be used if the alternate loop is temporary.

2.1.5 Availability and Reliability

A key to having a reliable system and maximizing availability is an adequate amount of redundant equipment to perform routine maintenance. If N represents the number of pieces to satisfy the normal cooling capacity, then often reliability standards are considered in terms of redundant pieces compared to the baseline of N. Some examples would be:

- $N + 1$—full capacity plus one additional piece
- $N + 2$—full capacity plus two additional pieces
- $2N$—twice the quantity of pieces required for full capacity
- $2(N + 1)$—full capacity plus one additional piece and the entire assembly repeated again (backup site)

A critical decision is whether N should represent just normal conditions or whether N includes full capacity during offline routine maintenance. For example, in an $N + 1$ redundancy scenario during routine maintenance of a single unit, the remaining capacity from N units is exactly what is required to meet the cooling load. If one of the N online units fails while another is under maintenance, the cooling capacity of the online units is insufficient.

The determination of the facility load represented by N should also be made based on local design conditions. ASHRAE has statistical design data based on 0.4%, 1%, and 2% excursions on an annual basis. The 0.4%, 1%, and 2% design conditions are exceeded 35, 87, and 175 hours, respectively, in a typical year. As an example, in Atlanta, the dry-bulb temperatures corresponding to the 0.4%, 1%, and 2% conditions are 93.9°F, 91.5°F, and 89.3°F, respectively (ASHRAE 2004a).

In an $N + 2$ redundancy scenario, even if one unit is down for maintenance, there is still an additional unit online above the required capacity that can compensate if one of the online units fails. This scenario means that there is always a redundant unit at all times, whereas $N + 1$ would have a redundant unit for the majority of the time but not during routine maintenance.

The $2N$ and $2(N + 1)$ systems can apply to pipe work as well as components. A $2N$ design will not need loop isolation valves since there is a complete backup system. However, due to the thermal inertia of the fluid in the systems, both systems must be operational for seamless transfer without disruption if one system detects a fault.

These configurations of $N + 1$ through $2(N + 1)$ exemplify the significant variation in the design of redundant systems to achieve equipment availability at least equal to N. Overlaid on top of these configurations is the influence of human error. Many failures are caused by human error; some configurations improve reliability from a mechanical standpoint but may increase the potential for human error.

Chilled-water or thermal storage may be applied to the data center central cooling plant to minimize or eliminate computer equipment shutdown in the event of a power failure to the datacom facility. Chilled-water storage tanks are placed in the building distribution chilled-water piping. In the event of a power failure, the chilled-water pumps use the UPS to provide a constant flow of chilled water to the equipment. The chilled-water tanks are used in conjunction with pumps to provide cooling water. The tanks should be sized to have enough storage capacity to match the battery run time of the UPS systems. Once the generators are operational, the chillers and pumps will resume normal operation.

For more information refer to *Design Considerations for Datacom Equipment Centers* (ASHRAE 2005a). The following chapters in that document provide helpful information regarding cooling services for equipment cooling systems:

- Chapter 4, "Computer Room Cooling Overview"—The relationship between methods of heat rejection as accomplished by direct expansion versus chilled-water systems is described, along with the functional characteristics and interrelationships of the refrigeration cycle, condensers, chillers, pumps, piping, and humidifiers. The chapter concludes with a description of control parameters and monitoring methodologies.
- Chapter 13, "Availability and Redundancy"—This chapter details aspects of availability, such as the concept of "five nines," failure prediction, mean time between failure (MTBF), and mean time to repair (MTTR). Concepts of redundancy, such as $N + 1$, $N + 2$, and $2N$ are introduced, defined, and discussed. Diversity and human error, as well as some practical examples of methods to increase availability and redundancy, are presented.
- Chapter 14, "Energy Efficiency"—Specific topics include chilled-water plants, CRAC units, fans, pumps, variable-frequency drives, humidity control, air- and water-side economizers, part-load operation, in-room airflow distribution, and datacom equipment energy efficiency.

2.2 EQUIPMENT

The facility cooling equipment is a broad topic and more detail can be obtained in the ASHRAE Handbooks. For the scope of this book, only the chillers, heat rejection equipment, energy recovery equipment, and pumps are covered.

2.2.1 Chillers

Chiller is the term used to describe mechanical refrigeration equipment that produces chilled water as a cooling output or end product (Figure 2.1). Basic information on chillers can be obtained from Chapter 38 of the *2004 ASHRAE Handbook—HVAC Systems and Equipment* (ASHRAE 2004a).

The chilled water produced by a chiller for building cooling is commonly 42°F to 45°F (5.5°C to 7°C) but can be selected and designed to produce temperatures that are higher or lower than this range. Often this temperature is too low for data center cooling. A warmer temperature will provide enough cooling without dehumidifying further and has the potential for significant energy savings. The downside to the higher temperature is that it takes more flow and, hence, larger pipes and/or pumps.

The chilled water can be 100% water or a mixture of water and glycol (to prevent the water from freezing if piping is run in an unconditioned or exterior area) plus other additives such as corrosion inhibitors. The capacity of chillers (and all other heat transfer coils, such as heat exchangers and cooling coils) is influenced or derated when glycol or additives are included.

In a typical chilled-water configuration, the chilled water is piped in a loop between the chiller and the equipment cooling system. It is important to consider the chiller's part-load efficiency, since data centers often operate at less than peak capacity.

One way to classify chillers is the basic method of chiller heat rejection (air-cooled or water-cooled). For air-cooled chillers, the compressors reject the heat they gain to the atmosphere using air-cooled condenser sections (Figures 2.2 and 2.3). The majority of air-cooled chillers are located outside to facilitate heat rejection to the atmosphere. However, due to spatial or other constraints, the air-cooled chiller components can be split from the heat rejection component (typically an air-cooled condenser) located remotely from the chiller.

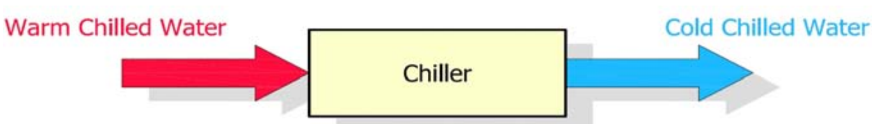

Figure 2.1 Generic chiller diagram.

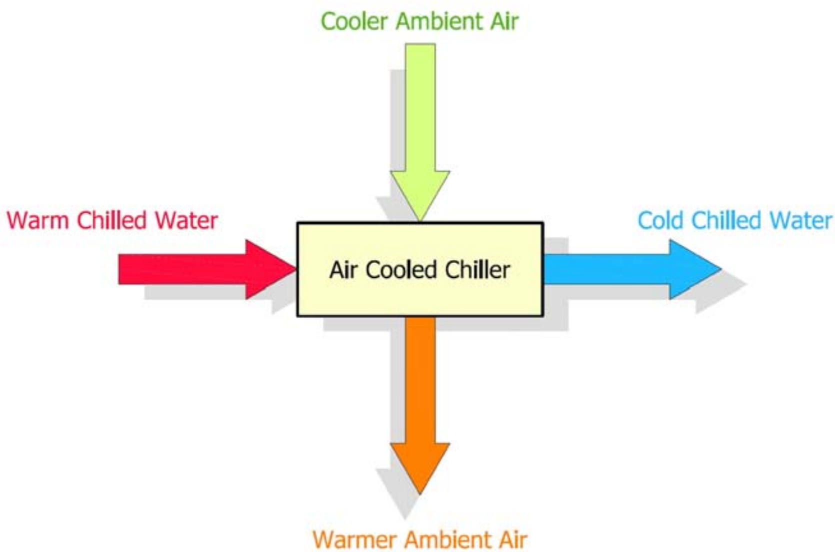

Figure 2.2 Schematic overview of a generic air-cooled chiller flow.

Figure 2.3 Typical packaged air-cooled chiller.

Water-cooled chillers use a second liquid loop called the *condenser water loop*. The condenser water loop is typically connected to an open or closed cooling tower (evaporative heat rejection unit) as shown in Figure 2.4.

A water-cooled chiller is shown in Figure 2.5. Based on electrical consumption per quantity of cooling produced and depending on the type of configuration used, water-cooled chillers are more energy efficient than the air-cooled equivalent. More information on chillers can be found in Chapter 38 of the *2004 ASHRAE Handbook—HVAC Systems and Equipment* (ASHRAE 2004a).

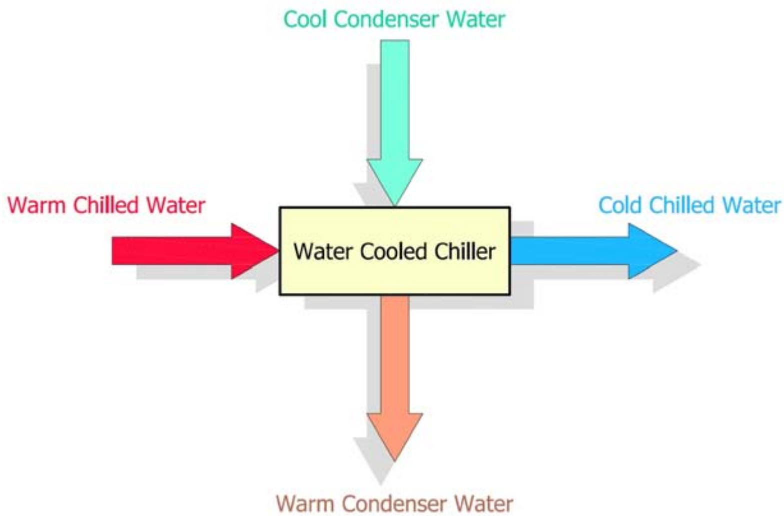

Figure 2.4 Schematic overview of a generic water-cooled chiller flow.

Figure 2.5 Water-cooled chiller.

2.3 HEAT REJECTION EQUIPMENT

A cooling tower is a heat rejection device that uses evaporative cooling (Figure 2.6). Cooling towers come in a variety of shapes, sizes, configurations, and cooling capacities. Since cooling towers require an ambient airflow path in and out, they are located outside, typically on a roof or elevated platform (Figure 2.7).

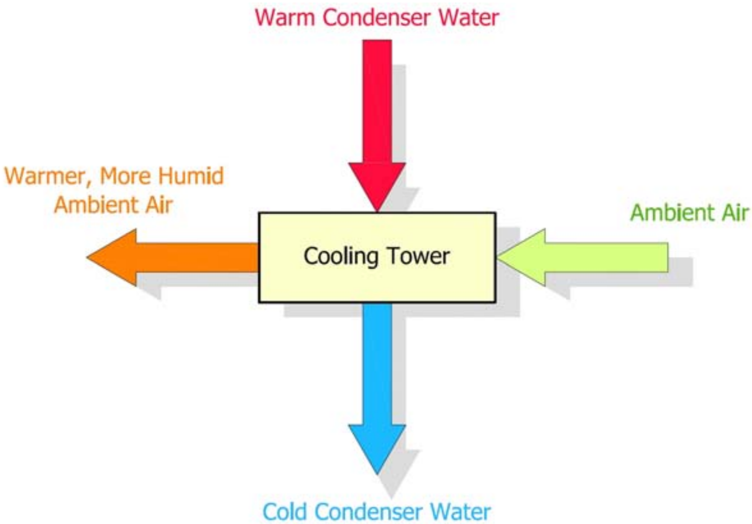

Figure 2.6 Schematic overview of a generic cooling tower flow.

Figure 2.7 Direct cooling towers on an elevated platform.

The elevation of the cooling tower relative to the remainder of the cooling system needs to be considered when designing the plant because the cooling tower operation and connectivity rely on the flow of fluid by gravity. Basic information on condenser water systems can be obtained from Chapter 13 of the *2004 ASHRAE Handbook—HVAC Systems and Equipment* (ASHRAE 2004a). Evaporative cooling needs a reliable water source. Depending on the quality of water and the extent of suspended solids that will stay in the remaining water after evaporation, a bleed rate to remove residue must be added to the evaporation rate. This is generally between 3.5 and 4 gallons per ton of refrigeration per hour but varies depending on location, elevation, and water chemistry. This will contribute to the total water storage or well source requirements.

The generic term *cooling tower* is used to describe both open-circuit (direct-contact) and closed-circuit (indirect-contact) heat-rejection equipment (see Figures 2.8 and 2.9). The indirect cooling tower is often referred to as a *closed circuit fluid cooler* or *fluid cooler.*

Where an open cooling tower is used, a heat exchanger should be considered to isolate the open water loop from the chilled water or other closed loop supplying the equipment cooling system to limit the possibility of fouling.

For data centers or other mission critical facilities, onsite water storage is a consideration. First, for water-cooled plants with evaporative cooling towers, makeup water storage could be provided to avoid the loss of cooling tower water following a disruption in water service. For large data centers, the tank size corresponding to 24 to 72 hours of onsite makeup water storage could be in the range of 100,000 to well over 1,000,000 gallons. Second, chilled-water storage could provide an emergency source of cooling. Placing the chiller plant on a UPS can be very costly; chilled-water storage may offer a short period of cooling to keep chillers off a UPS. Prolonged power outages, however, will still require that chillers be fed from emergency generators to maintain operations.

Typically the selection of cooling towers is an iterative process since there are a number of variables resulting in the opportunity for trade-offs and optimization. Some of those variables include size and clearance constraints, climate, required operating conditions, acoustics, drift, water usage, energy usage, part-load performance, etc.

2.4 PUMPS

Pumps and pumping system design should take into account energy efficiency, reliability, and redundancy. It may be possible to design pumps with variable-speed drives so that the redundant pump is always operational. Ramp-up to full speed occurs on a loss of a given pump. Use of multiple pumps running at slower speeds will also save energy. Premium efficiency motors should always be specified since the payback period will be short with 24/7 operation.

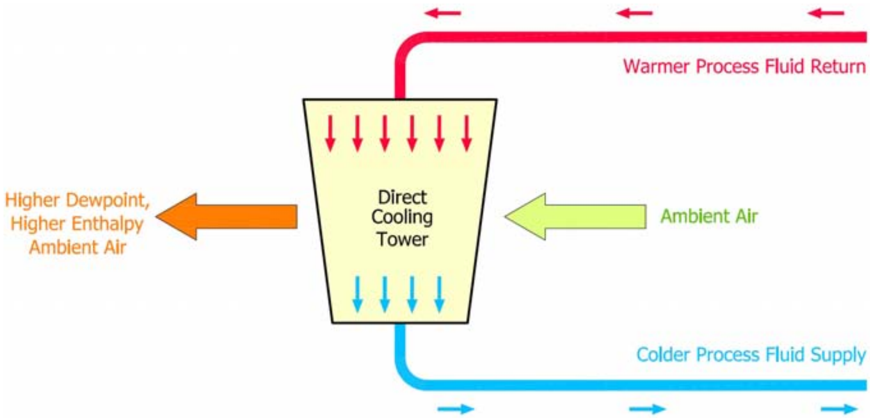

Figure 2.8 Direct or open circuit cooling tower schematic flow diagram.

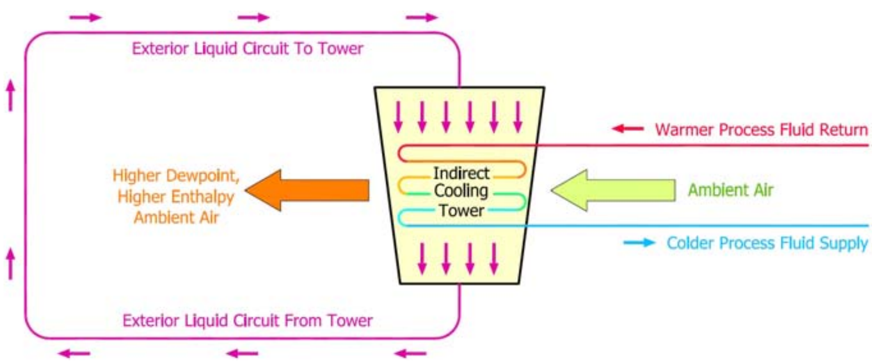

Figure 2.9 Indirect cooling tower schematic flow diagram.

Basic information on pumps can be obtained from Chapter 39 of the *2004 ASHRAE Handbook—HVAC Systems and Equipment* (ASHRAE 2004a). Piping systems are covered in the same volume of the Handbook in Chapter 12.

2.5 ENERGY RECOVERY EQUIPMENT

Fundamentally, the economizer process involves utilizing favorable ambient weather conditions to reduce the energy consumed by the facility cooling systems. (The IT industry still has some concerns regarding air-side economizers, but these

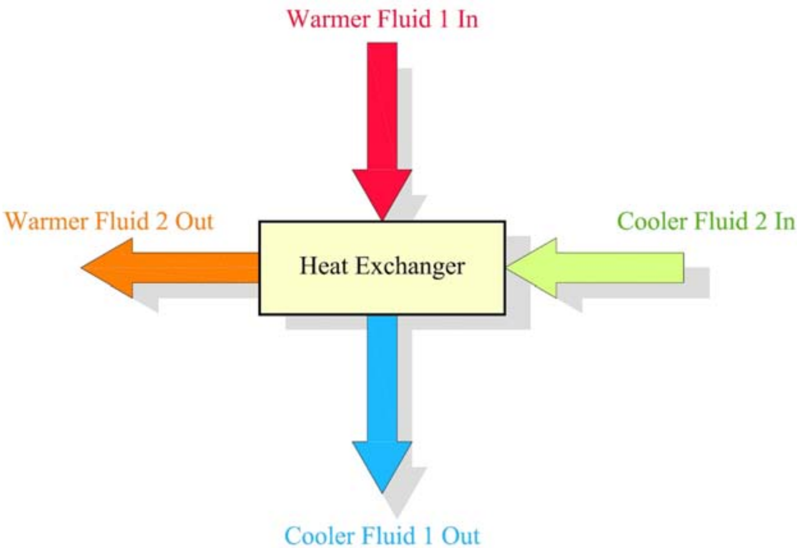

Figure 2.10 Simple overview of the heat exchanger process.

are beyond the scope of this liquid cooling document.) Most often this is accomplished by limiting the amount of energy used by the mechanical cooling (refrigeration) equipment. Since the use of an economizer mode (or sequence) reduces the energy consumption while maintaining the design conditions inside the space, another term that is used in the building cooling industry for economizer mode operation is *free cooling*.

ANSI/ASHRAE Standard 90.1-2004, Energy Standard for Buildings Except Low-Rise Residential Buildings, is a "standard" or "model code" that describes the minimum energy efficiency standards that are to be used for all new commercial buildings and include aspects such as the building envelope and cooling equipment. This standard is often adopted by an authority having jurisdiction (AHJ) as the code for a particular locale.

Water-side economizer cycles often use heat exchangers to transfer heat from the condenser-water system to the chilled-water system as conditions permit. As its name suggests, a heat exchanger is a device that relies on the thermal transfer from one input fluid to another input fluid. One of the fluids that enters the heat exchanger is cooler than the other entering fluid. The cooler input fluid leaves the heat exchanger warmer than it entered, and the warmer input fluid leaves cooler than it entered. Figure 2.10 is a simplistic representation of the process. Figure 2.11 shows a typical installation of plate and frame heat exchangers.

Figure 2.11 Insulated plate and frame heat exchangers.

3

Facility Piping Design

3.1 GENERAL

The piping architecture defines the relationship between the cooling source (plant) and the load (electronic equipment). The architecture should consider simplicity, cost, ease of maintenance, ease of upgrade/change, ease of operation, controls, reliability, energy usage, etc.

Typically the basic options for piping architecture are established and then reviewed for their effectiveness within the spatial constraints of the facility. For example, a loop may look like a good option from a piping architecture perspective but the routing paths available may not provide the space and locations needed to create effective loops.

This chapter is organized into:

- Spatial considerations, including routing
- Basic piping architecture
- Piping arrangements for the central plant
- Water treatment issues
- Earthquake protection

One other important consideration is pipe sizing criteria. Analysis of plant, distribution, and terminal pipe sizes results in trade-offs between capital and operational costs. Larger pipe sizes yield lower water velocities, which, in turn, lower pumping power (smaller pumps) and operational costs. Generally, velocities should be as high as practical without sacrificing system integrity. (See discussion of velocity limits in Section 5.1.1.3). Typically, increased pumping energy does not outweigh the lower capital cost or the space savings associated with the smaller pipe sizes. Pipe sizing also affects how much additional cooling load can be added to the data center at a future date. As cooling and power densities continue to grow, data centers must be scalable to house this future growth. One strategy is to oversize the chilled-water piping plant mains and distribution headers to accommodate future load increases. Oversizing these will save energy and allow smaller pumps for much of the data center's life.

The chilled-water piping must be planned and designed so the data center can remain operational while adding more computer equipment.

3.2 SPATIAL CONSIDERATIONS

Usually the spatial considerations start by determining what piping can be located in the data center and where. Some examples of stakeholder concerns/preferences are:

- No overhead piping, minimal overhead piping, or no constraint on overhead piping other than to avoid routing directly above electronic equipment.
- All piping mains and pipes above a certain size can be run in the data center or are confined to pipe galleries, chases, troughs, utility pits, etc.

The spatial constraints do not just apply to routing but also to the accessibility of valves and terminations as well as the impact of piping penetrations through firewalls. Stakeholder piping location preferences combined with the physical constraints of the facility often have as much or more influence on the piping architecture as any other influence, such as operations, energy, or redundancy. Consideration should be given to laying out large piping first so as to minimize pressure drop, e.g., large radius bends, few changes of direction, etc. Then arrange pumps, chillers, etc., to tie into the piping with minimal changes of direction. This will likely result in chillers and pumps arranged in less than a linear fashion—perhaps at 45° instead of 90°.

Although each layout and design is different, a good starting point is to allocate cross-sectional dimensions that are twice the pipe diameter. Some examples:

- If the pipe is 6 in. in diameter, then allow 2 × 6 in. or 12 × 12 in. in cross section.
- If the pipe is 18 in. in diameter, then allow 2 × 18 in. or 36 × 36 in. in cross section.

These allocations are to allow for the pipe flanges, pipe guides/supports, valve handles, etc. Also, the piping design must consider expansion and contraction as well as seismic considerations. For example, the piping may be stored and installed in an ambient temperature of 90°F (32°C) and operate at 45°F (7°C). The differential temperature (90°F – 45°F = 45°F [25°C]) can cause significant movement during system startup and cooldown.

3.3 BASIC PIPING ARCHITECTURE

Inherent to the concept of liquid cooling is extending the infrastructure that contains the cooling media to an area local to the datacom equipment. In addition, the infrastructure delivery method is typically at a much finer granularity (i.e., many more points of connection) based on the terminal use of the liquid. Cooling equipment may be serving a few racks, a single rack, or perhaps even multiple components within a single rack.

The infrastructure itself is typically a copper or steel piping network. Many different piping architectures and flow principles can be used to extend the piping network to the rack and the increased points of connection while at the same time providing an overall system that is consistent with the necessary reliability and flexibility of a datacom facility. The following group of diagrams within this section represent some of the different piping architectures and flow principles that can be used.

Direct Return (Figure 3.1)

A direct return system is the most basic type of piping system and is used in traditional HVAC design where there are a reduced number of connection points. In this system, the supply and return piping is fed in a radial manner and the loads that are closest to the cooling plant have the shortest supply piping lengths *and* the shortest return piping lengths.

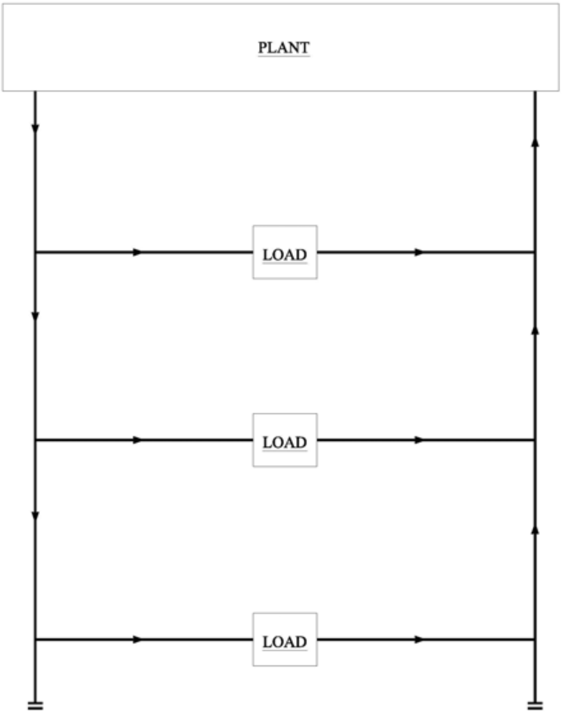

Figure 3.1 Example of direct return flow principle.

However, the direct return method, when used in an application that has many close connection points, may require an excessive amount of balancing valves to ensure proper system operation. This is due to the variation in supply and return piping lengths to a given load.

Advantages

1. Least expensive to construct, uses a minimal amount of pipe, valves, and fittings.
2. Simplest to operate and understand.

Disadvantages

1. Least reliable since only one source of cooling exists.
2. No redundancy in piping to the load. Any pipe failure or leak or future addition could jeopardize system availability.
3. May require additional balancing valves.

Reverse Return (Figure 3.2)

The objective of the reverse return flow principle is to inherently create a piping network with an element of self-balancing. This is achieved by having the loads supplied by piping closest to the cooling plant also be the loads that are at the most remote end of the return piping and vice versa. This is achieved by essentially having the flow in the return piping parallel the flow in the supply piping as it feeds the various loads around the building. This results in the combined length of supply and return piping for any given load being approximately equal, which creates a system that can be considered self-balancing.

Advantages

1. Simple to operate and understand.
2. Self-balancing.

Disadvantages

1. Less reliable; again, only one source of cooling.
2. No redundancy in pipe or chilled-water routes. Routine maintenance or system expansion could require complete system shutdown.
3. A little more expensive to install than direct return (i.e., more piping required).

Looped Mains Piping Schemes

The remaining piping architecture examples illustrated in this section involve the use of looped piping mains. Looped piping mains involve a closed loop that is tapped at various points to feed loads. The flow of liquid within the loop can occur in two directions from the source and, in theory, there is a “no-flow zone” near the midpoint of the loop.

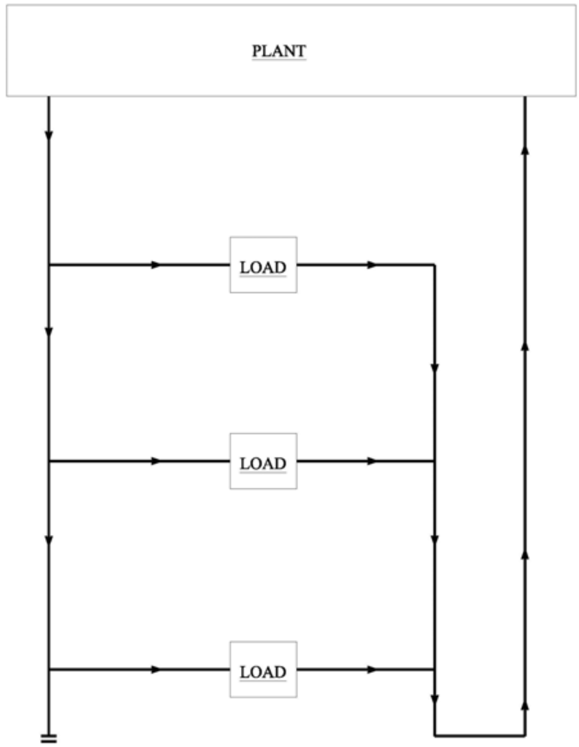

Figure 3.2 Example of reverse return flow principle.

More importantly, the architecture of looped piping mains also allows for a section of main piping to be isolated for maintenance or repair. Loads that were downstream of the isolated section can then be backfed from the other side of the looped mains to allow for greater online availability of the cooling system.

Various piping architectures can be used to create a loop design. These variations allow a system to attain different levels of isolation, different hydraulic characteristics, and different levels of modularity. The following diagrams illustrate some examples of looped piping architectures.

Single-Ended Loop with Direct Feed (Figure 3.3)

A single-ended loop has a single point of connection (supply and return piping) to the plant. The piping is typically looped within the datacom area and, in this particular configuration, the loads are directly fed from the mains' loop piping.

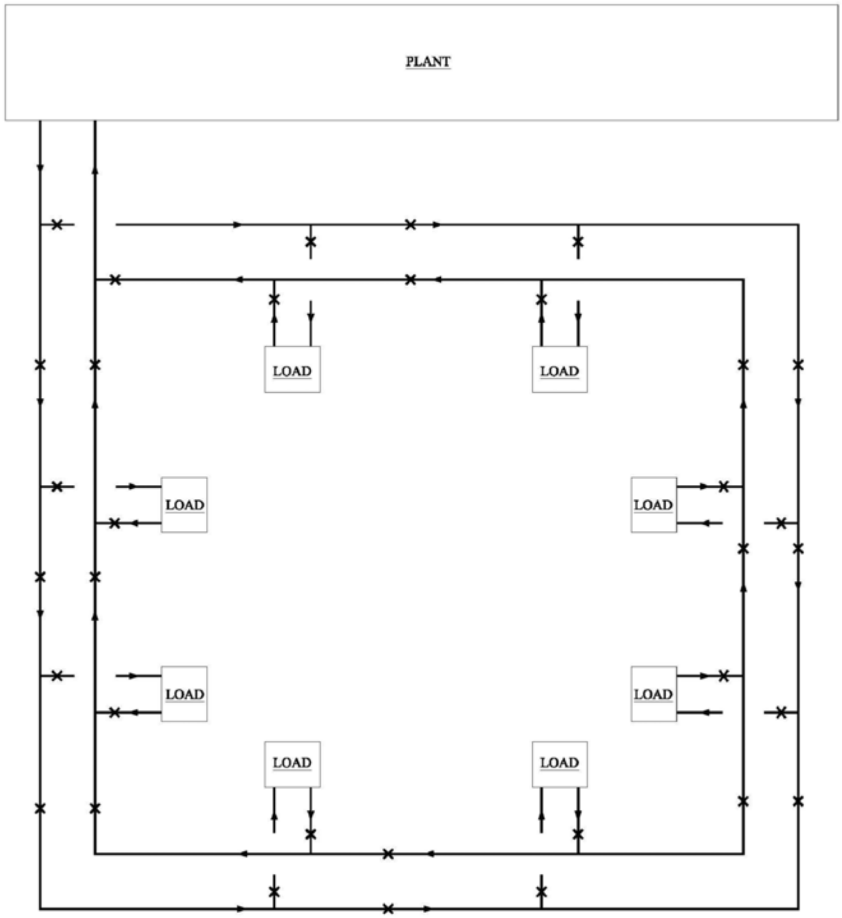

Figure 3.3 Single-ended loop with direct feed.

A popular application of this piping architecture is for an air-cooled CRAC unit based cooling system, where CRAC units are located around the perimeter of the datacom area.

Advantages

1. Self-balancing.
2. Increased reliability over direct and reverse returns systems with two piping routes to the load.

3. Individual pipe sections and future equipment installations are serviceable without system shutdown.

Disadvantages

1. Increased complexity and understanding.
2. Increased installation costs.

Single-Ended Loop with Common Cross Branches (Figure 3.4)

Similar to the previous example (Figure 3.3), the same single-ended loop architecture is used. The difference is in the connection of the loads, which are now indirectly fed from cross-branch piping connected at two locations to the mains' loop. The cross-branch piping is said to be *common* since it is used by multiple loads on both sides of its route as a supply and return flow path.

This method not only allows for a bidirectional flow of liquid in the mains but also within each cross branch. As such, it provides multiple paths for flow to reach the majority of the loads should a section of the mains' loop or the cross branch need to be isolated for reasons of maintenance or repair.

Advantages

1. Increased reliability with multiple piping routes to load.
2. Self-balancing.
3. Used primarily for water-cooled rack units.
4. Individual pipe sections and future equipment installations are serviceable without system shutdown.

Disadvantages

1. Increased installation costs.
2. Increased operational complexity.

Single-Ended Loop with Dedicated Cross Branches (Figure 3.5)

The same single-ended loop is used here as with the previous two examples (Figures 3.3 and 3.4) and the same cross-branch piping as in the previous example (Figure 3.4).

Now, however, the indirect connections of the loads are supplied from an increased number of cross-branch pipes. This allows for an increase in granularity of the loads and, therefore, an increased level of reliability (i.e., the isolation of a section of cross-branch piping will not impact as many loads since there are fewer connections per cross branch).

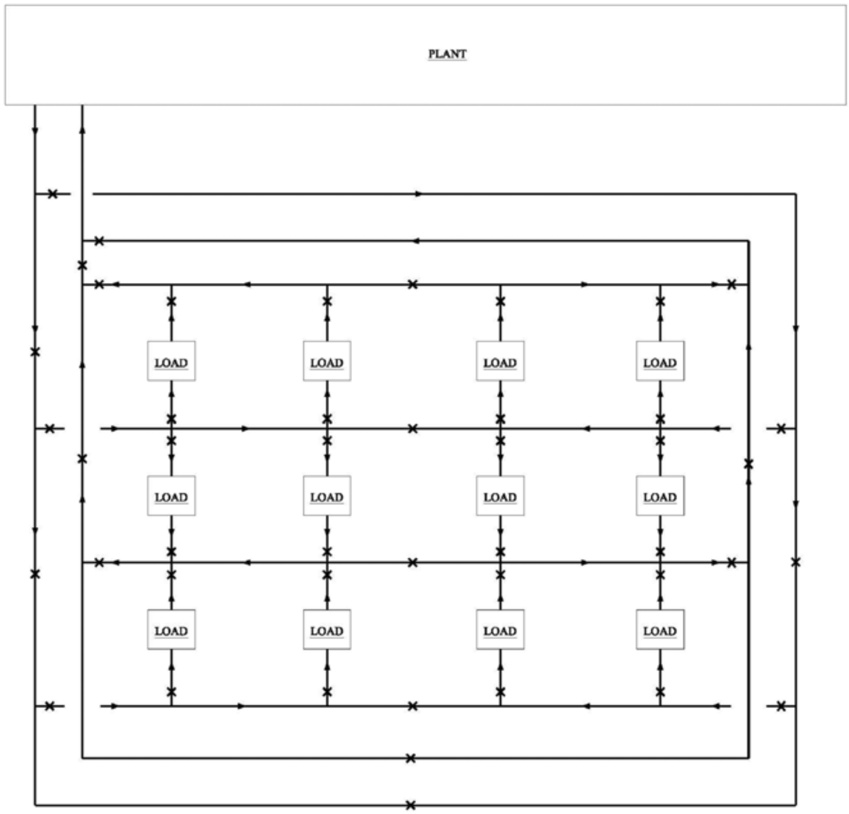

Figure 3.4 Single-ended loop with common cross branches.

As is apparent from the diagram, this architecture involves many more cross branches, so the increased granularity needs to be evaluated against increased cost and hydraulic/control complexity.

Advantages

1. Increased reliability with multiple piping routes to load.
2. Self-balancing.
3. Individual pipe sections and future equipment installations are serviceable without system shutdown.

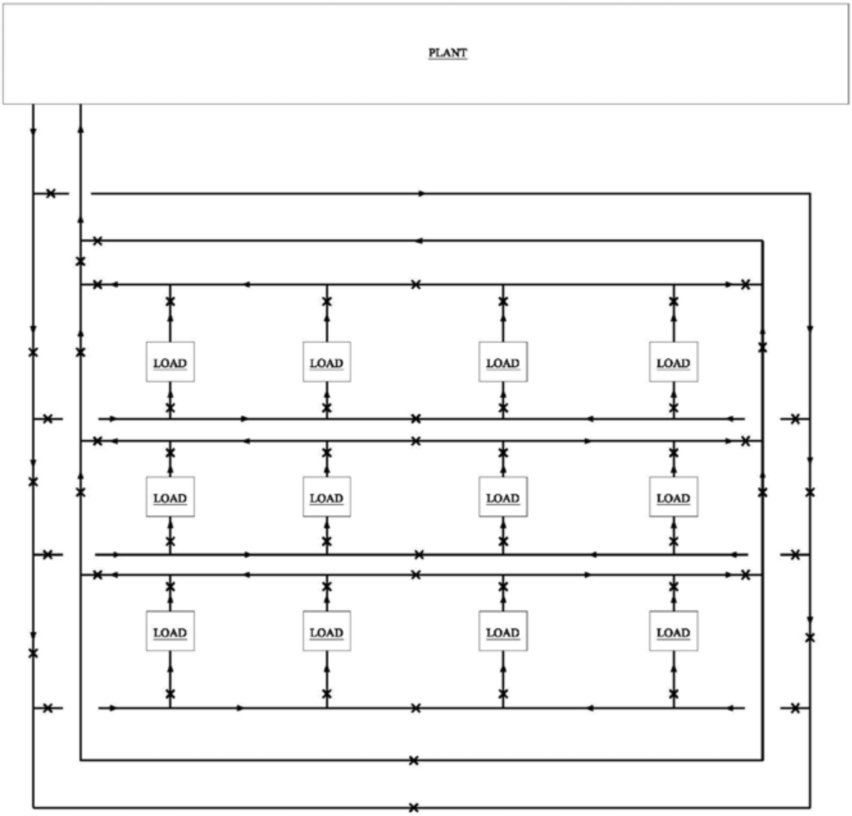

Figure 3.5 Single-ended loop with dedicated cross branches.

Disadvantages

1. Increased installation costs.
2. Increased operational complexity.

Double-Ended Loop with Direct Feed (Figure 3.6)

The only difference between the single-ended loop (shown in Figure 3.3) and the double-ended loop (shown in Figure 3.6) is that in this piping architecture there are two connections to the plant, which eliminates the single point of failure that exists for all single-ended loop piping configurations (e.g., if a need exists to isolate

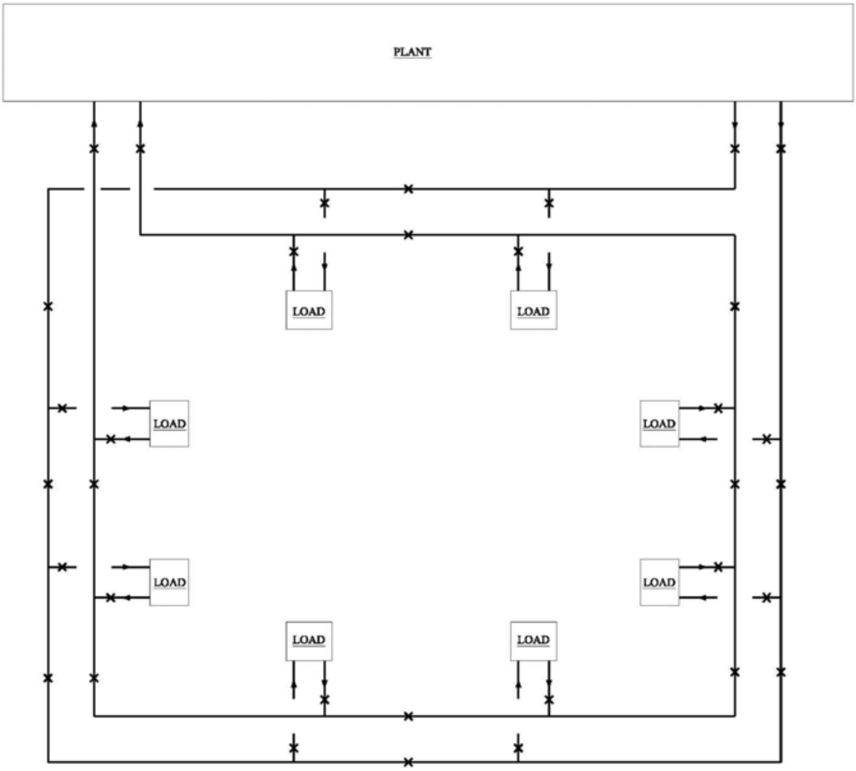

Figure 3.6 Double-ended loop with direct feed.

the piping between the connection to the plant and upstream toward the plant itself, this method will still allow cooling to all loads via the second connection).

To illustrate an even greater level of reliability, consider the second connection to the loop to be from an operationally independent plant. These independent plants could even be in geographically different locations within the same facility.

Advantages

1. High reliability.
2. Redundant piping routes to load and a second cooling supply and return mains from the plant.
3. Redundant cooling supply and return piping from a second central plant.

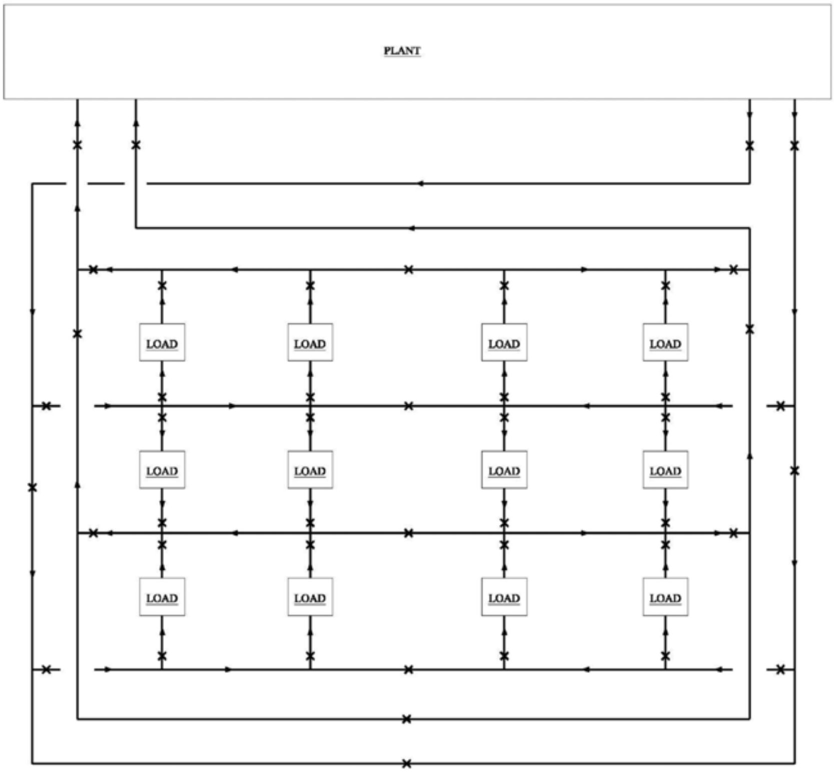

Figure 3.7 Double-ended loop with common cross branches.

4. Individual pipe sections and future equipment installations are serviceable without system shutdown.
5. Self-balancing.

Disadvantages

1. Increased installation costs.
2. Increased operational complexity.

Double-Ended Loop with Common Cross Branches (Figure 3.7)

As stated previously, the difference between this piping architecture and the single-ended loop with common cross branches (shown in Figure 3.4) is that two connections to the plant are made to eliminate the single point of failure.

Advantages

1. High reliability.
2. Redundant piping routes to load and a second cooling supply and return mains from the plant.
3. Redundant cooling supply and return piping from a second central plant.
4. Individual pipe sections and future equipment installations are serviceable without system shutdown.
5. Self-balancing.

Disadvantages

1. Increased installation costs.
2. Increased operational complexity.

Double-Ended Loop with Dedicated Cross Branches (Figure 3.8)

Similarly, the principal difference between this piping architecture and the single-ended loop with dedicated cross branches (shown in Figure 3.5) is that two connections to the plant are made to eliminate the single point of failure.

Advantages

1. High reliability.
2. Redundant piping routes to load and a second cooling supply and return mains from the plant.
3. Redundant cooling supply and return piping from a second central plant.
4. Individual pipe sections and future equipment installations are serviceable without system shutdown.
5. Self-balancing.

Disadvantages

1. Increased installation costs.
2. Increased operational complexity.

3.4 PIPING ARRANGEMENTS FOR THE COOLING PLANT

The cooling plant equipment can be configured various ways. For example, the chillers can be configured in series or parallel and have different preferential loading schemes. The pumping and flow can be configured as constant flow, stepped variable flow, or variable flow. The building owner or occupant will have to perform an engineering analysis to determine which configuration is best for their data center.

Figure 3.9 shows a typical decoupled or condenser-water system/chilled-water system pumping configuration.

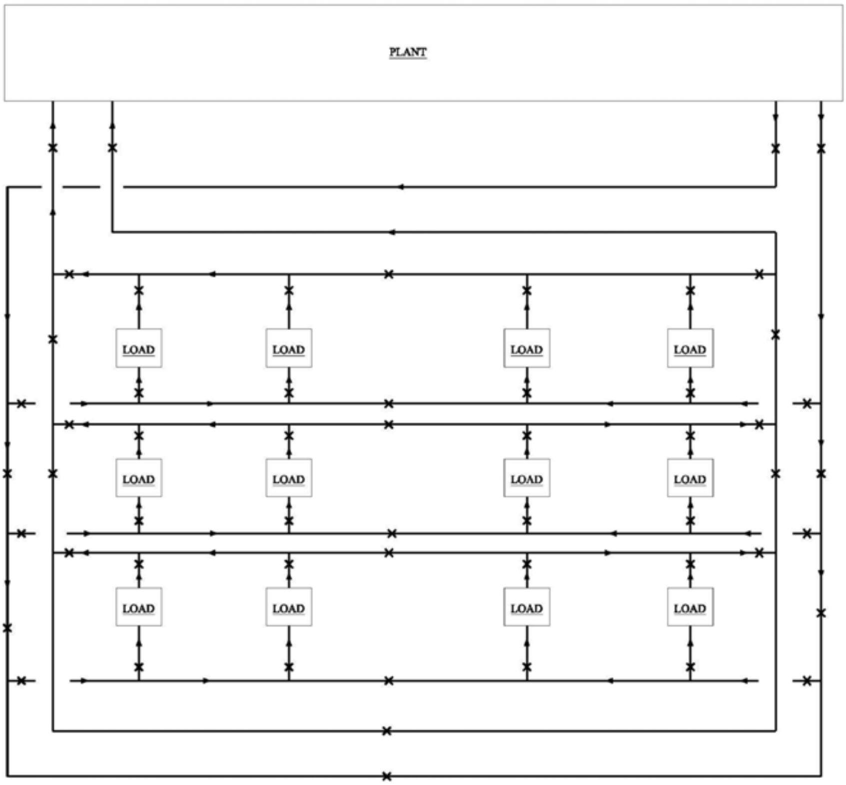

Figure 3.8 Double-ended loop with dedicated cross branches.

3.4.1 CHWS Pipe Sizing

The pipe conveyance has to support real load rather than an average bulk load on the raised floor. Today's data center has an eclectic grouping of various loads ranging from cross-connect racks (no-load) to blade servers. The CHWS pipe work must provide cooling at hot spot areas during the life of the data center. Since these areas might change with new equipment over time, flexibility to provide hotter areas with cooling must be oversized in the CHWS distribution. Example: if the average load is at a design density of 100 W/ft^2, the CHWS distribution should be able to supply any local area of the raised floor with 175–200 W locally. In today's environment of changing technology, all CHWS piping should plan for a series of additional water taps off the distribution to serve the future requirements for auxiliary cooling equipment.

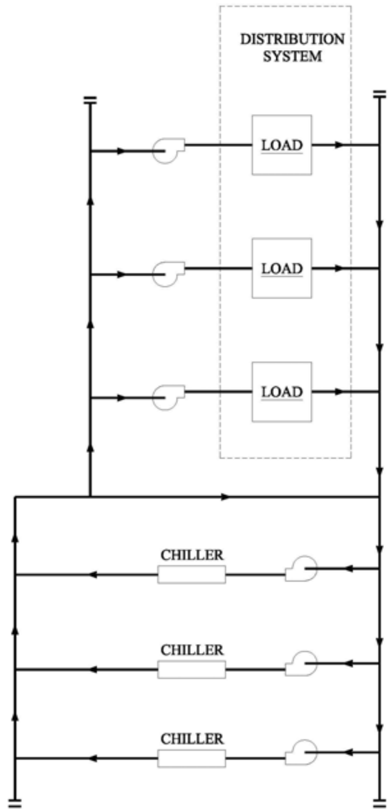

Figure 3.9 Condenser water system/chilled-water system distribution piping.

3.4.2 Loop Isolation Valve Failures

The basis for reliable service in the mechanical plant and pipe work uses loop isolation valves to isolate the operating system from the repair or expansion. If the isolation valve itself fails, both sides of the loop are then exposed to the next isolation valves. Depending on the design configuration, this might include the redundant pump or chiller and render the capacity of the system below critical load. Therefore, the entire loop strategy relies on the loop isolation valve reliability and service life. To remove this single point of failure from the system, either a double-pipe system (independent of the first system) or a double set of isolation valves is used in locations where a single valve can affect critical load.

3.5 WATER TREATMENT ISSUES

Whenever an oversized pipe system with redundant pathways and expansion sections is prebuilt for future service and then installed into a data center, the operation of the center needs water treatment to protect the system. Pipe walls interact with the water they circulate. Water treatment includes testing and injecting chemicals to prevent the pipe walls from corrosion. Pipe loops may not have any flow at all for extended periods, and expansion sections may never have flow. The life of the treatment chemicals is approximately 30 days (depending on strength and water quality). Where pipe flow is low or is not disturbed, fouling and corrosion may occur. Periodically the flow in these pipes must be accelerated to full flow for a short time each month to re-protect the pipe walls. In stagnant segments, the water flow must be directed to the areas to re-coat those surfaces. Allowing the operations staff to gain experience shifting cooling water conveyance is also important for operational readiness in the event of a leak. The expansion legs should have a circulation jumper (1/4 in. bypass) with valves to allow a flushing of these pipes monthly to keep protection and fouling under control. The entire process can be automated if loop isolation and VFD drives can be controlled by the automation system.

3.6 EARTHQUAKE PROTECTION

Depending on the design area and the risk aversion of the data center insurance underwriters, the mechanical and fire protection pipes over 4 in. in diameter may come under bracing requirements for earthquake protection. A qualified structural engineer will be required to design for these requirements.

4

Liquid Cooling Implementation for Datacom Equipment

Chapter 1 briefly describes several examples of different implementations of liquid cooling technologies. More importantly, Chapter 1 clearly defines air-cooled and liquid-cooled datacom equipment, as well as air-cooled and liquid-cooled racks or cabinets. The following sections provide detailed descriptions of the various liquid cooling implementations. Section 4.1 illustrates cooling at the rack/cabinet level, showing several common configurations for air, liquid, and combination (air + liquid) cooling. A more detailed look inside the rack/cabinet is provided in Section 4.2, which discusses air- and liquid-cooled datacom equipment. Because liquid cooling often employs a coolant distribution unit (CDU) to condition and circulate the coolant, Section 4.3 provides a detailed overview of this system. The words *fluid* and *coolant* are used interchangeably throughout. Throughout this chapter, the liquid lines are shown in the figures as a solid or dashed line and airflow is represented with arrows that contain dots.

4.1 OVERVIEW OF LIQUID-COOLED RACKS AND CABINETS

A rack or cabinet is considered to be liquid-cooled if liquid must be circulated to and from the rack or cabinet for operation. The following figures illustrate cooling at the rack/cabinet level. The first is a basic air-cooled rack. The remaining figures show other options that utilize liquid cooling or a combination of air cooling and liquid cooling. The figures in this section all show the coolant supply and return lines under the raised floor. Other facility implementations may allow such lines to be routed above the floor or from the ceiling. Coolant supply and return connections for the rack/cabinet can be from the base, side, or top.

Figure 4.1 shows a purely air-cooled rack or cabinet implementation. While this book's focus is on liquid cooling, this figure provides a baseline with which the vast majority of datacom center operators are familiar. It should be noted that while the figures all show a front-to-back configuration for the airflow and the rack, it can also be configured as front-to-top or front-to-back-and-top (ASHRAE 2004b).

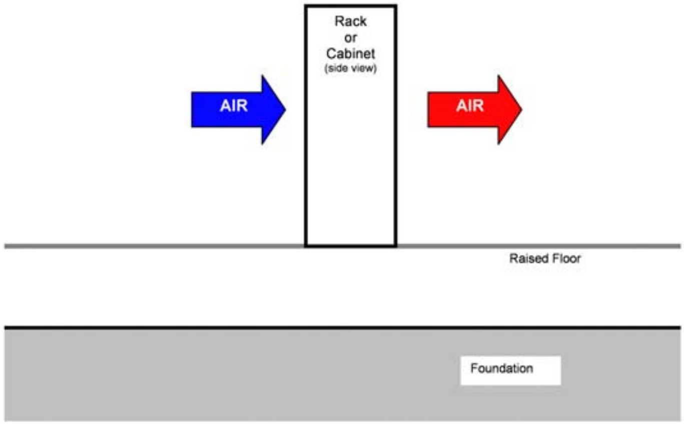

Figure 4.1 Air-cooled rack or cabinet.

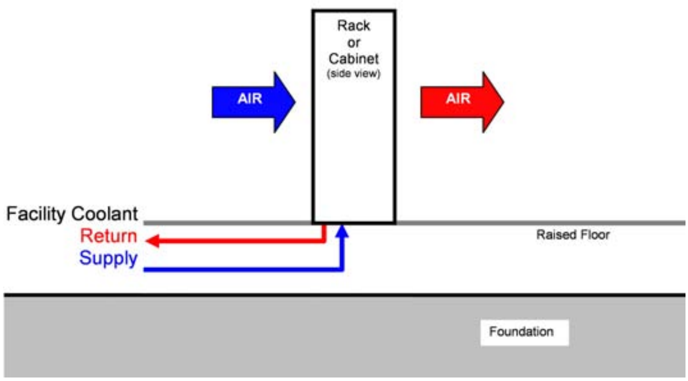

Figure 4.2 Combination air- and liquid-cooled rack or cabinet.

Figure 4.2 shows a combination air-cooled and liquid-cooled rack or cabinet that could receive the chilled working fluid directly from some point within the CHWS or CWS loop. By the definitions in Chapter 1, the rack or cabinet is liquid-cooled since coolant crosses the interface between the facility and the rack or cabinet. One implementation could have the electronics air-cooled, with the coolant removing a large percentage of the waste heat via a rear door heat exchanger. Another implementation

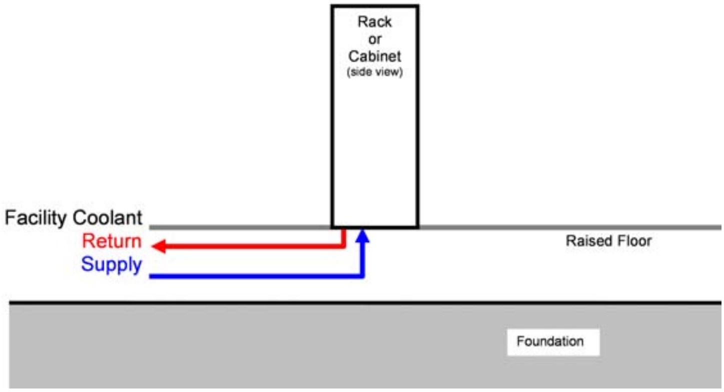

Figure 4.3 Liquid-cooled rack or cabinet (side view).

could have the coolant delivered to processor spot coolers (some form of cold plate), with the balance of the electronics being air-cooled. The descriptions provided are two of many different implementations and should not be understood as the only possible implementations. Further details of the implementation within the rack, cabinet, or datacom equipment are provided in Section 4.2. It is important to note that this configuration is susceptible to condensation because there is no CDU in place to raise the temperature of the chilled fluid above dew point, if necessary.

Figure 4.3 shows a purely liquid-cooled rack or cabinet. One example of such an implementation may have all the electronics in the rack or cabinet conduction-cooled via cold plates. This cooling method could deploy water, refrigerant, or other dielectric coolant as the working fluid. Another implementation may have all the electronics cooled via liquid flow-through (e.g., forced flow boiling), jet impingement, spray cooling, or another method that deploys a dielectric coolant to directly cool the electronics. Yet another implementation would include a totally enclosed rack that uses air as the working fluid and an air-to-liquid heat exchanger. Further details are provided in Section 4.2. Similar to Figure 4.2, the configuration in Figure 4.3 is also susceptible to condensation because there is no CDU in place to raise the temperature of the chilled fluid above dew point, if necessary.

Figure 4.4 shows a combination air-cooled and liquid-cooled rack or cabinet with an external CDU. The CDU, as the name implies, conditions the technology cooling system (TCS) or datacom equipment cooling system (DECS) coolant in a variety of manners and circulates it through the TCS or DECS loop to the rack, cabinet, or datacom equipment. This implementation is similar to that of Figure 4.2, with the exception that there is now a CDU between the facility (CHWS or CWS) level supply of chilled fluid and the rack or cabinet. This implementation allows the CDU

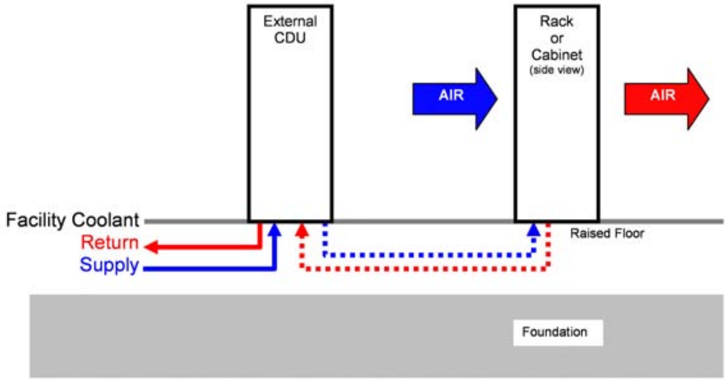

Figure 4.4 Combination air- and liquid-cooled rack or cabinet with external CDU.

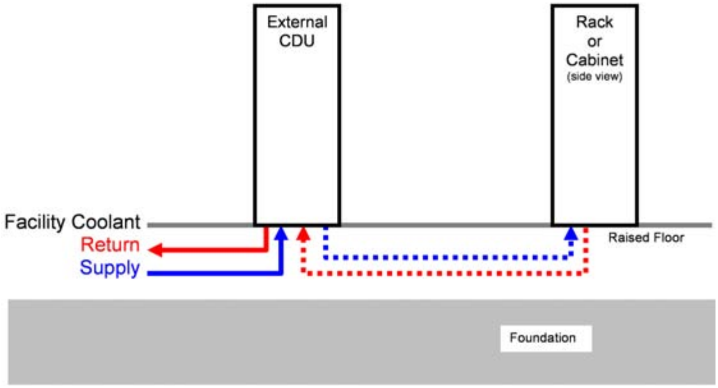

Figure 4.5 Liquid-cooled rack or cabinet with external CDU.

to condition the coolant delivered to the rack or cabinet to a temperature above the facility's dew point.

Figure 4.5 shows a purely liquid-cooled rack or cabinet implementation. This implementation is similar to that of Figure 4.3, as well as Figure 4.4, where an external CDU is included.

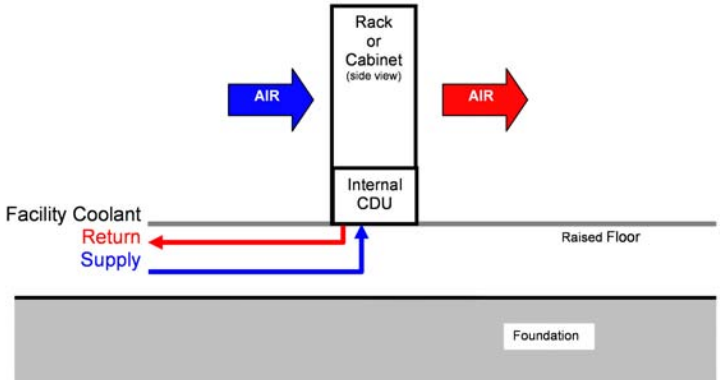

Figure 4.6 Combination air- and liquid-cooled rack or cabinet with internal CDU.

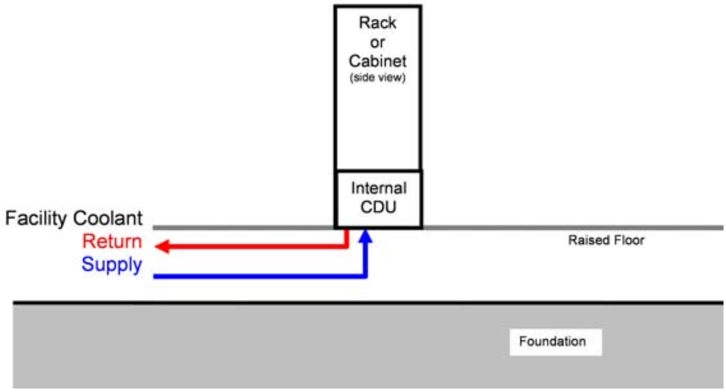

Figure 4.7 Liquid-cooled rack or cabinet with internal CDU.

Figures 4.6 and 4.7 are the final implementations to be discussed in this section. These implementations have a lot in common with the implementations of Figures 4.4 and 4.5, respectively. One obvious difference is the fact that the racks or cabinets shown in Figures 4.6 and 4.7 now possess dedicated CDUs, i.e., internal CDUs. The CDUs are shown at the bottom of the rack, but other configurations could include them on the side or top of the rack. This implementation provides more flexibility to the datacom center operator in that the racks or cabinets can now condition their coolants to vastly different conditions as a function of the workload or the electronics within. Another benefit is that different coolants (e.g., water, refrigerant, dielectric) can now be deployed in the different racks as a function of workload or electronics type. Additional detail is provided in Section 4.2.

4.2 OVERVIEW OF AIR- AND LIQUID-COOLED DATACOM EQUIPMENT

While Section 4.1 dealt with the external interaction between the building infrastructure and the rack or cabinet, this section examines the possible liquid cooling systems internal to the rack or cabinet. Described within this section are seven datacom cooling configurations, which represent the most commonly used configurations. Other configurations are possible but are not examined here for brevity's sake.

The systems remove heat from the rack or cabinet by means of a fluid whose chemical, physical, and thermal properties are conditioned by a CDU for that purpose. The conditioned fluid may be water, an antifreeze mixture, dielectric fluid, or refrigerant.

Figure 4.8 shows a rack or cabinet with combined air and liquid cooling. The computer room air and the conditioned liquid each remove a part of the heat load. In this configuration, air is the only coolant entering the datacom equipment. The air-to-liquid heat exchanger extracts heat from the air that is either entering or leaving the datacom equipment and, as a result, reduces the heat load on the computer room air conditioning and reduces hot air recirculation. The heat exchanger may be mounted on the rack or cabinet doors or in any other location along the airstream.

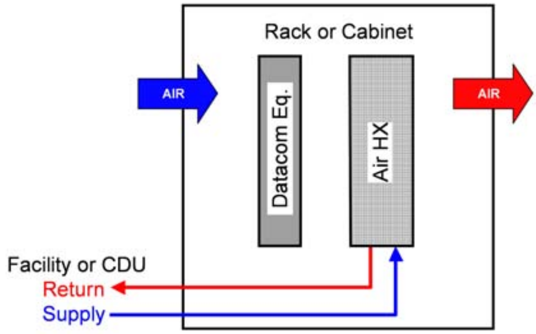

Figure 4.8 Open air-cooled datacom equipment in an air/liquid-cooled rack.

Whichever the case, care should be taken to consider heat exchanger effectiveness and condensation, especially with configurations where the heat exchanger is upstream from the datacom equipment. Finally, if the air pressure drop across the heat exchanger is sufficiently low, no additional fans will be required to achieve the necessary airflow rate.

Figure 4.9 shows an enclosed cabinet where air is again the only coolant entering the datacom equipment, but here computer room air is excluded, and the entire heat load (apart from heat loss though panels) is removed by the conditioned fluid. Additional fans will probably be required to drive sufficient airflow through the system. To cool the datacom equipment, this configuration requires the same amount of cooling air as the configuration shown in Figure 4.1. The fans and the heat exchanger may be located in various positions. For example, the fans may be in a duct on the rear door, and the heat exchanger may be at the bottom of the cabinet. The objectives are to achieve sufficient airflow to extract the total heat load and to achieve an airflow distribution that avoids uneven cooling and "hot spots" within the cabinet.

Figure 4.10 illustrates a rack or cabinet where the internal heat transfer is by liquid only. The heat exchanger and pump are shown as "optional" to recognize the case where the facility working fluid from the CHWS or CWS loop is sufficiently conditioned to allow direct entry to the datacom units. Unless there is a local CDU associated with the cabinets, one important benefit of the CDU, namely, the limiting of fluid leakage volume, is lost.

Figure 4.11 shows the addition of computer room air cooling to a Figure 4.10 system. In view of the variety of heat sources and their different form, we could expect this to be a more frequent configuration since airflow through complex or heterogeneous geometries is easier to implement with air than with liquids.

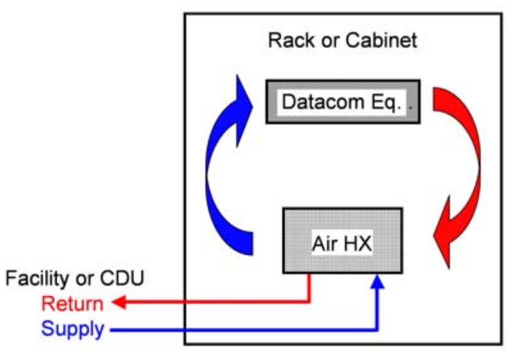

Figure 4.9 Closed air-cooled datacom equipment in a liquid-cooled cabinet.

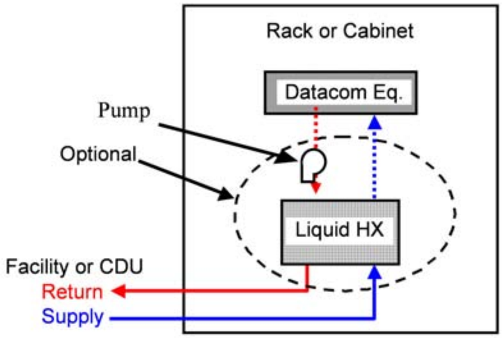

Figure 4.10 Liquid-cooled datacom equipment in a liquid-cooled rack.

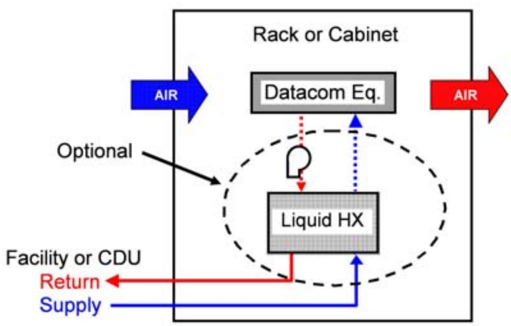

Figure 4.11 Open air- and liquid-cooled datacom equipment in an air/liquid-cooled rack.

The arrangement in Figure 4.12 goes one step further. Like Figure 4.11, it shows a combination of air and liquid cooling, but now the air is a closed loop that stays within the cabinet. The air is cooled by a separate heat exchanger. This configuration may represent a cabinet housing multiple pieces of datacom equipment, some liquid-cooled and others air-cooled. It may also represent datacom equipment with some components that are liquid-cooled (e.g., the CPU) and other components that are air-cooled (e.g., the power supplies).

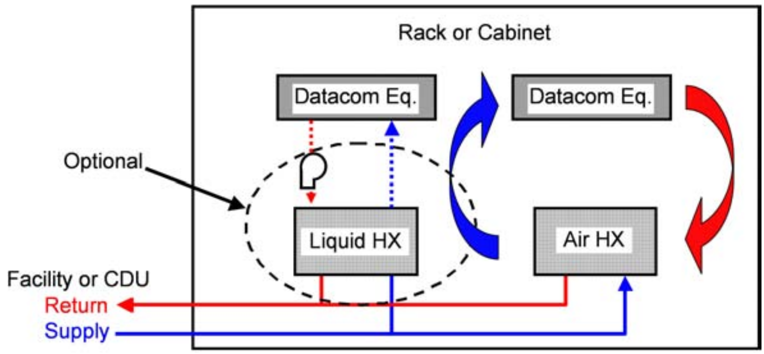

Figure 4.12 Closed air- and liquid-cooled datacom equipment in a liquid-cooled rack.

Figures 4.13 and 4.14 illustrate the application of vapor-compression cycle heat transfer within the rack or cabinet. Figure 4.13 employs an internal DECS loop, and in Figure 4.14 air is the working fluid between the datacom equipment and the vapor-compression system.

In all of the systems listed above, the principal heat "sink" is the facility conditioned fluid in the CHWS (chiller) or CWS loop, in some cases assisted by computer room conditioned air. When using a DX system, the facility conditioned fluid would be liquid refrigerant. When significant heat loads are removed by conditioned fluids, the computer room air temperature may be allowed to rise, presenting a more comfortable operating environment. However, the current focus on fluids is driven primarily by escalating heat dissipations and the subsequent risk of thermal failure and shortened life span that are associated with high chip operating temperatures, as well as lower energy costs because liquid solutions are more efficient than air.

4.3 OVERVIEW OF COOLANT DISTRIBUTION UNIT (CDU)

The coolant distribution unit (CDU), as the name implies, conditions the TCS or DECS coolant in a variety of manners and circulates it through the TCS or DECS loop to the rack, cabinet, or datacom equipment (see Figure 1.1 for cooling loop descriptions). For example, the CDU can condition the coolant for temperature and cleanliness (particulate and chemical). The CDU implementations shown in Figures 4.15–4.20 all demonstrate coolant temperature control via the heat exchanger. It should be noted that coolant temperature control can also be achieved via heat exchanger bypass. The CDU can also be designed to condition the coolant if the TCS or DECS coolant is a refrigerant or dielectric. The rack or facility coolant supplied to the CDU can be refrigerant

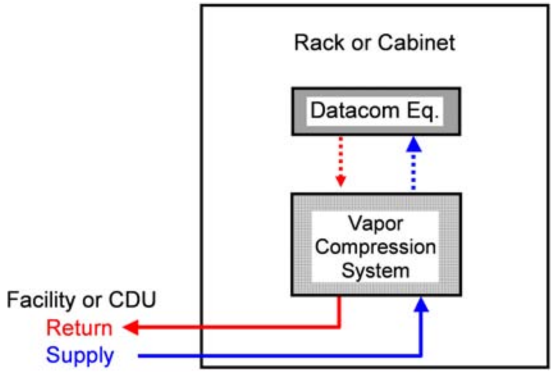

Figure 4.13 Liquid-cooled datacom equipment in a liquid-cooled rack using a vapor-compression system.

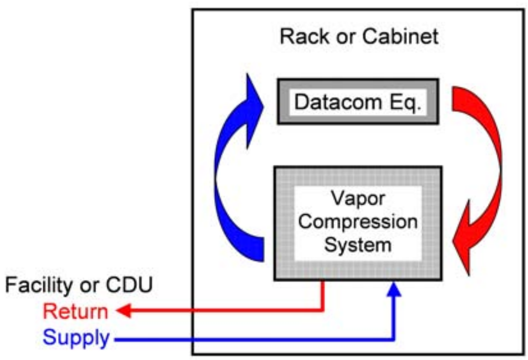

Figure 4.14 Air-cooled datacom equipment in a liquid-cooled cabinet using a vapor-compression cycle.

supplied by a DX system. The functionality of the CDU will depend, in large part, upon the coolant that is being conditioned and the level at which the CDU is implemented (i.e., internal to the datacom equipment, internal to the rack/cabinet, or at facility level and external to the rack). In all cases, the CDU includes a pump or compressor to circulate the fluid through the coolant loop. While all of the figures show counterflow heat exchangers, other types are possible.

The CDU provides a number of important benefits, as follows:

- **Prevention of condensation:** It provides an opportunity for temperature conditioning, which could allow the coolant to be delivered to the electronics above the dew point.
- **Isolation:** It may allow the electronics to be isolated from the harsher facility water in the CHWS or CWS loop. The loop supplied by the CDU also utilizes a lower volume of coolant, so a coolant leak is less catastrophic.
- **Coolant flexibility:** The separate loop associated with the CDU allows users the flexibility to use any number of coolants.
- **Temperature control:** The separate loop associated with the CDU allows users the flexibility of running the electronics at a desired temperature. This temperature can be above, at, or below the dew point.

Figures 4.15–4.20 illustrate the breadth of current implementations of CDUs. While not shown, it is implied that some form of quick disconnect (or quick coupling and a ball valve) is implemented at all points where coolant crosses the "CDU boundary" line, which is indicated by a dashed-dotted line around the CDU. The use of quick disconnects allows for low loss of coolant and facilitates replacement of a CDU in the event of a catastrophic failure. It is important to keep in mind that there are many possible implementations of CDUs outside of those shown in Figures 4.15–4.20. Figures 4.15 and 4.16 illustrate implementations internal to the datacom equipment. Figures 4.17 and 4.18 illustrate implementations internal to the rack/cabinet. Figures 4.19 and 4.20 illustrate implementations external to the rack/ cabinet and at the facility level. Further descriptions follow.

Figure 4.15 shows an implementation of a CDU internal to the datacom equipment. The datacom equipment rejects the heat from its electronics via a cold plate that docks with another rack/cabinet-based cold plate. At a very basic level, the CDU consists of a cold plate, an accumulator (or a reservoir), and a pump (or multiple pumps for redundancy). The operation of the pump is regulated by some form of controller, which is in communication with one or more sensors within the CDU. The data from the sensors are used to control the level of operation of the CDU.

Similarly to Figure 4.15, Figure 4.16 illustrates an internal datacom equipment implementation of a CDU. The basic functionality for this implementation is similar to that for Figure 4.15, with the exception of the use of a liquid-to-liquid heat exchanger to reject the datacom equipment heat. In other implementations, the CDU could deploy a liquid-to-air heat exchanger.

Figure 4.17 shows a CDU implementation internal to the rack or cabinet, as shown in Figures 4.6 and 4.7. The CDU shown consists of a liquid-to-liquid heat exchanger, an accumulator (or a reservoir), a pump (or multiple pumps for redundancy), a chemical bypass filter for solvent and water removal, and a particulate filter. Applications of an internal CDU will often use refrigerant or dielectric as the

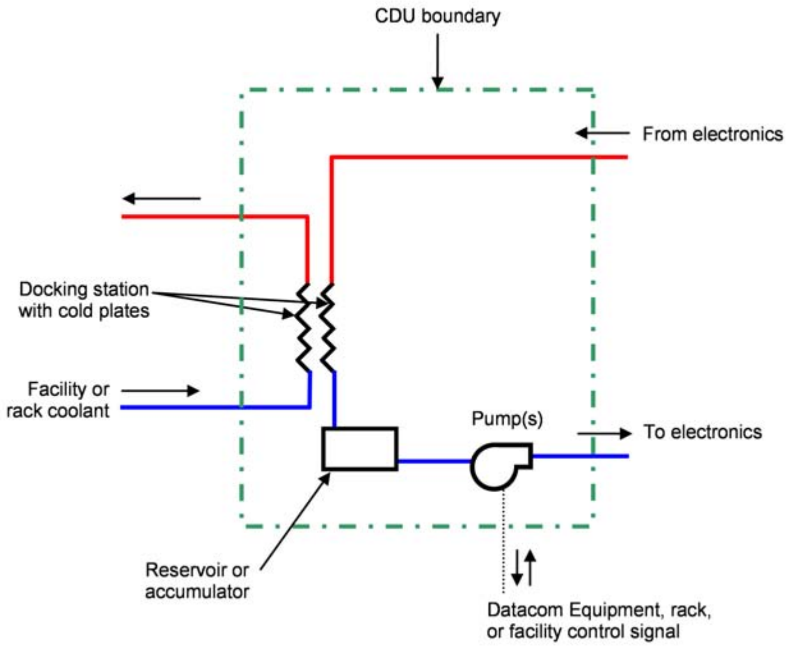

Figure 4.15 Internal datacom equipment-based CDU that utilizes a docking station and cold plates.

cooling fluid. These fluids necessitate the use of additional filters. This implementation also interfaces with one or more controllers that control the operation of the CDU. The controllers can use sensor inputs at the datacom equipment, rack, or facilities level. While a liquid-to-liquid heat exchanger is shown, this implementation could also deploy a liquid-to-air heat exchanger.

Similarly to Figure 4.17, Figure 4.18 illustrates the internal implementation of a CDU at the rack/cabinet level, as shown in Figures 4.6 and 4.7. The basic functionality of this CDU is similar to that shown in Figure 4.17, with the main difference being the use of a vapor-compression system. A key benefit of a vapor-compression system is that it allows the user to drop the temperature of the working coolant below that of the coolant to which it is rejecting the heat from the electronics. Unique to the vapor-compression system, a compressor (or multiple ones for redundancy) and an expansion valve are used. While a liquid-cooled condenser is shown, an air-cooled condenser could also be deployed.

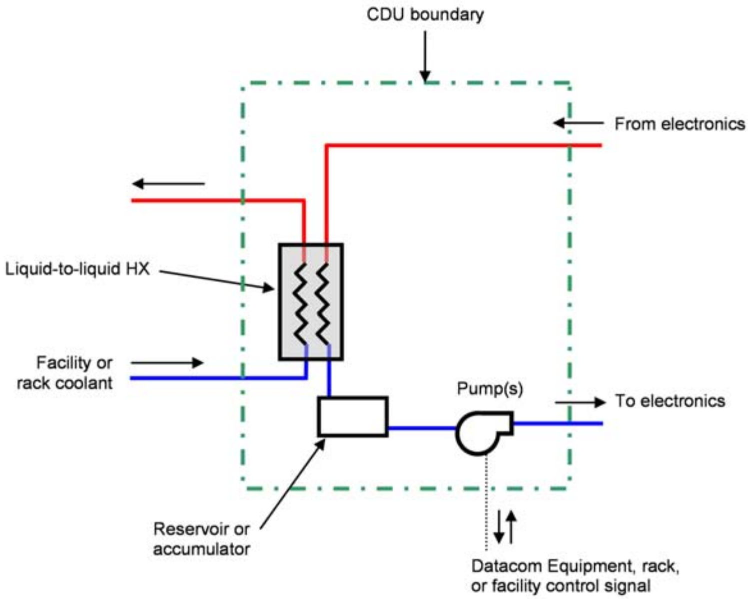

Figure 4.16 Internal datacom equipment-based CDU that utilizes a liquid-to-liquid heat exchanger.

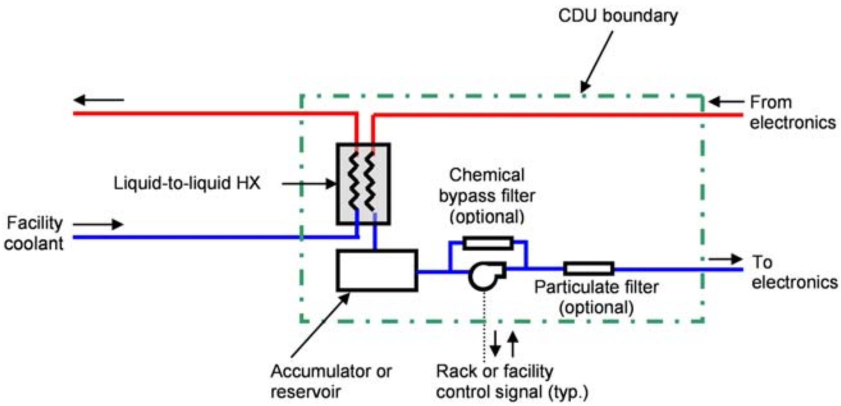

Figure 4.17 Internal rack or cabinet-based CDU that utilizes a liquid-to-liquid heat exchanger.

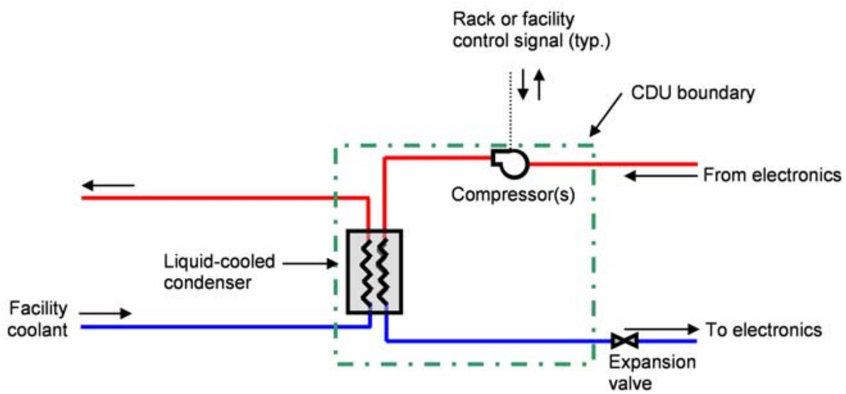

Figure 4.18 Internal rack or cabinet-based CDU that utilizes a liquid-cooled condenser and a vapor-compression system.

Figure 4.19 shows a CDU implementation external to the rack and at the facility level, as shown in Figures 4.4 and 4.5. This form of CDU can supply cooling fluid to a single rack or a row of racks. In its most basic form, this CDU consists of a liquid-cooled condenser, a pump (the additional space typically allows for redundant pumps), and an expansion tank. The CDU communicates with one or more controllers at the datacom equipment, rack/cabinet, and facility levels, which can control the operation of the CDU. It is possible, but probably not likely, that an air-cooled condenser could be deployed. This is probably not likely due to the amount of heat that has to be rejected. One of the key benefits of this type of implementation is that it isolates the racks/cabinets and datacom equipment from the facility water in the CHWS or CWS loop, which tends to be much harsher on the internal materials.

Figure 4.20 shows the implementation of a CDU external to the rack or at the facility level, as shown in Figures 4.4 and 4.5. The basic functionality is similar to that of Figure 4.19, with the key difference being the use of a vapor-compression system. Similar to Figure 4.19, the basic vapor-compression system uses a compressor (or multiple ones for redundancy), a liquid-cooled condenser, and an expansion valve. This implementation of the CDU can also communicate with one or more controllers at the datacom equipment, rack/cabinet, and facility levels. Similar to the implementation of Figure 4.20, it is possible, but probably not likely, that an air-cooled condenser could be deployed. This implementation also isolates the racks/cabinets and datacom equipment from the facility water in the CHWS or CWS loop.

The temperature in the datacom equipment can rise very, very quickly if the CDU fails. Therefore, CDU implementations need to be redundant and fault-toler-

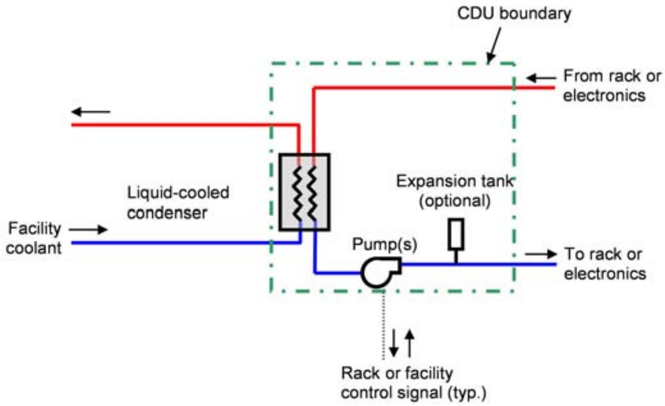

Figure 4.19 Facility-based CDU that utilizes a liquid-cooled condenser.

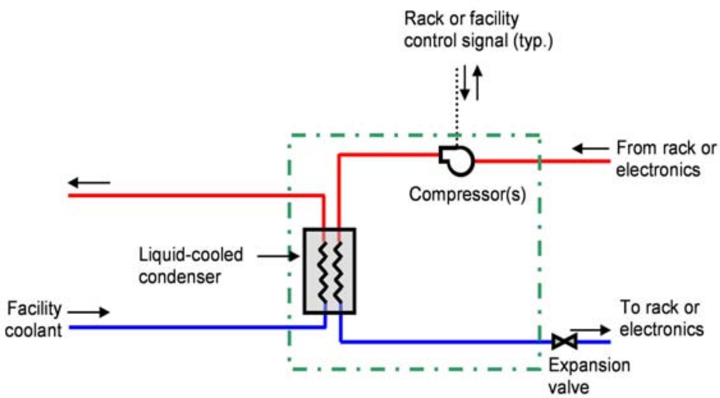

Figure 4.20 Facility-based CDU that utilizes a liquid-cooled condenser and a vapor-compression system.

ant. All of the strategies discussed in Section 2.1.5 regarding CRAC availability and redundancy apply to the CDU. Multiple CDUs should be implemented to eliminate downtime from mission critical installations. The chiller supplying the cooling fluid to the CDU also needs to be redundant and fault-tolerant to ensure it is always supplying cooling water to the CDU.

The CDU vendor also must design the CDU to be fault-tolerant. For instance, the coolant pump in the CDU is most likely the point of failure. Most systems will have redundant pumps. The pump will either alternate on a relatively short-term basis (weekly or monthly) or one pump will be the primary pump and a secondary pump will be tested periodically to make sure it is operational. The pumps will likely have isolation valves so that the nonfunctional pump can be isolated for replacement without bringing the entire CDU (and associated datacom equipment) down for repair. On smaller systems that may not accommodate dual pumps, the vendor will select a pump with a long life. In this situation, it is expected that the CDU will be upgraded before the pump wears out. Power to the CDU also needs to be backed up by a UPS. This will allow the datacom equipment to continue running or to shut down gracefully in the event of a power outage.

5

Liquid Cooling Infrastructure Requirements for Chilled-Water Systems

This section defines interface requirements between the facilities controlled by the building facility operators and the datacom equipment controlled by the datacom manufacturers. Demarcation lines for the CHWS are provided to describe where these interfaces occur in relationship to the datacom equipment and its location within the data center or telecom room.

5.1 BUILDING FACILITY CHILLED-WATER SYSTEMS (CHWS)

Figures 5.1 and 5.2 show the interfaces for an external liquid-cooled rack with remote heat rejection. The interface is located at the boundary at the facility CHWS loop and does not impact the TCS or DECS loops, which will be controlled and managed by the cooling equipment and datacom manufacturers. However, the definition of the interface at the CHWS loop affects both the datacom equipment manufacturers and the facility operator where the datacom equipment is housed. For that reason all of the parameters that are key to this interface will be described in detail in this section. Also, note that the coolant loops in the CDU are typically separate.

5.1.1 Chilled-Water Details

The majority of liquid cooling applications are anticipated to be water or water plus additives (e.g., propylene glycol, etc.). The following sections focus on these applications. Different design and operational requirements will certainly be required for alternate liquids, e.g., refrigerants or dielectrics (see Section 5.2 for non-chilled-water systems). Datacom equipment that requires a supply of chilled water from the data center facilities will typically adhere to the operational requirements contained within this section.

5.1.1.1 Operational Requirements. The datacom equipment should generally accommodate chilled-water supply temperatures that range from 40°F to 60°F (4°C to 16°C). The recommended nominal is 50°F (10°C) water. The actual CHWS loop

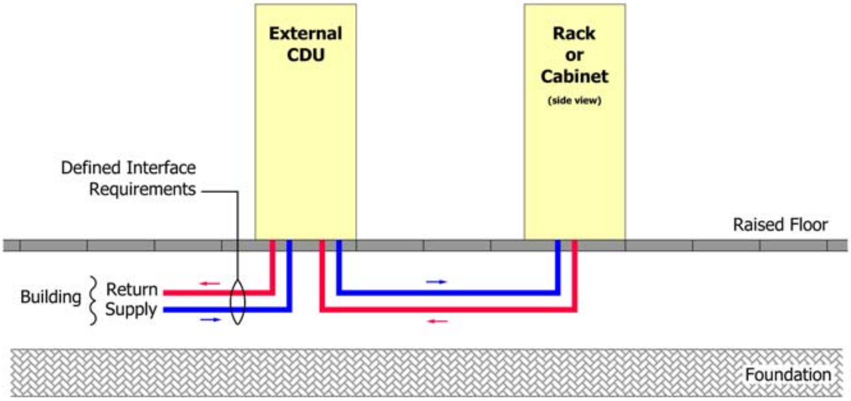

Figure 5.1 Combination air- and liquid-cooled rack or cabinet with external CDU (same as Figure 4.5).

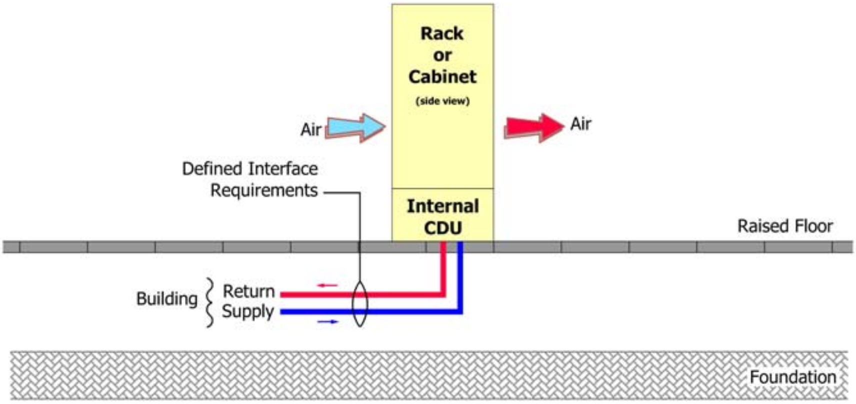

Figure 5.2 Combination air- and liquid-cooled rack or cabinet with internal CDU (same as Figure 4.6).

temperature may be set by a campus-wide operational requirement. It also may be the optimum of a balance between lower operational cost using higher-temperature chilled-water systems versus a lower capital cost with low-temperature chilled-water systems. Consideration of condensation prevention is a must. In the CHWS loop, insulation will typically be required. In the TCS and DECS loops, condensation control is typically provided by an operational temperature above the dew point.

The chilled-water supply temperature measured at the inlet of the datacom equipment or the CDU should not exceed a rate of change of 5°F (3°C) per five-minute cycle. This may require that the infrastructure be on UPS electrical supply.

The maximum allowable water pressure supplied to the TCS and DECS loops should be 100 psig (690 kPa) or less.

The chilled-water flow rate requirements and pressure-drop values of the datacom equipment vary, depending on the chilled-water supply temperature and the percentage of treatment (antifreeze, corrosion inhibitors, etc.) in the water. Manufacturers will typically provide configuration-specific flow rate and pressure drop values that are based on a given chilled-water supply temperature (nominally 50°F [10°C]).

The CHWS should incorporate a water flow control valve, typically a two-way modulating water flow valve. Generally, these valves are open and modulating when the datacom equipment is operating and fully closed when the datacom equipment is powered off. Depending on the design of the chilled-water system, a method of bypassing chilled water or de-energizing chilled-water pumps when the datacom equipment is not operating may be needed. Alternatively, for systems that must maintain a constant flow rate, three-way valves can provide this same function.

The flow rate for chilled water to be delivered to the datacom equipment or the CDU can be determined from flow rate tables provided by the manufacturer of such equipment.

5.1.1.2 Chilled-Water Flow Rates. Chilled-water flow rates are shown in Figure 5.3 for given heat loads and given temperature differences. Temperature differences typically fall between 9°F and 18°F (5°C and 10°C).

Fouling factors refer to an additional thermal resistance that occurs when the heat transfer surfaces have deposits built up over a period of time. These deposits are present in the coolant water. The higher the fouling factor becomes, the higher the required chilled-water flow rate becomes.

Typical fouling factors derived from $m^2 °C/W$ are:

Muddy or silty	0.0006
Hard (over 15 grains/gal)	0.0006
City water	0.0002
Treated cooling tower	0.0002

Refer to the list of references on how to apply these factors in sizing heat exchangers.

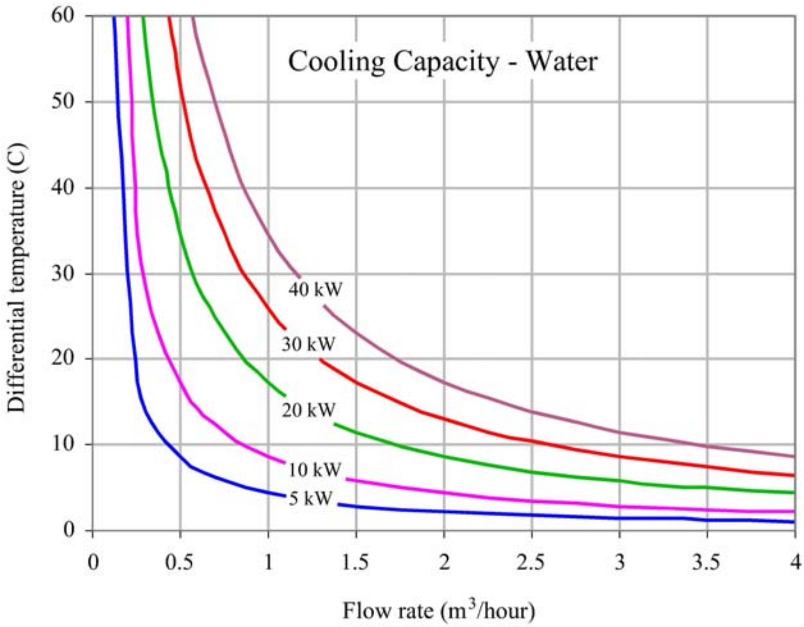

Figure 5.3 Typical chiller water flow rates for constant heat load.

5.1.1.3. Velocity Considerations. The velocity of the water in the CHWS loop piping must be controlled to ensure mechanical integrity is maintained over the life of the system. Velocities that are too high can lead to erosion, sound/vibration, water hammer, and air entrainment. Particulate free water will have less water velocity damage to the tubes and associated hardware. Table 5.1 provides guidance on maximum chilled-water piping velocities for systems that operate over 8,000 hours per year. Flexible tubing velocities should be maintained below 5 ft/s (1.5 m/s).

5.1.1.4. Liquid Quality/Composition. Table 5.2 identifies the water quality requirements that are necessary to operate the liquid-cooled system. The reader is encouraged to reference Chapter 48 of the *2003 ASHRAE Handbook—HVAC Applications*. This chapter, titled "Water Treatment," provides a more in-depth discussion about the mechanisms and chemistries involved.

5.1.1.4.1 Water Quality Problems. The most common problems in cooling systems are the result of one or more of the following causes:

- **Corrosion: the dissolution of metals by water.** Corrosion can be in the form of general corrosion (entire surface) or localized corrosion (pitting or stress),

Table 5.1 Maximum Velocity Requirements

Pipe Size	Maximum Velocity (fps)	Maximum Velocity (m/s)
>3 in.	7	2.1
1.5 to 3 in.	6	1.8
<1 in.	5	1.5
All flexible tubing	5	1.5

Table 5.2 Water Quality Specifications for the Chilled-Water System (CHWS) Loop

Parameter	Recommended Limits
pH	7 to 9
Corrosion inhibitor	Required
Sulfides	<10 ppm
Sulfate	<100 ppm
Chloride	<50 ppm
Bacteria	< 1000 CFUs/mL
Total hardness (as $CaCO_3$)	<200 ppm
Residue after evaporation	<500 ppm
Turbidity	<20 NTU (Nephelometric)

which causes metal perforation and rapid failure. The primary metals that can corrode in the loop are aluminum, stainless steel, and copper. Corrosion of steel and copper is often of a general nature, although the majority of system failures result from localized corrosion and not general corrosion. Corrosion of aluminum is often in the form of pitting. Typically, pH values below neutral (pH < 7) add significantly to the corrosion of typical metal piping systems. Corrosion is also driven by high levels of chlorides, sulfides, and sulphates in the loop.

- **Fouling: insoluble particulate matter in water.** Insoluble particulate matter settles as a result of low flow velocity or adheres to hot or slime-covered surfaces and results in heat-insulating deposits and higher pressure drops in the loop. A deposit is generally iron with small amounts of copper and mineral scales such as calcium carbonate and silt. Fouling is related to the amount of particulate matter or total suspended solids in the fluid. A full loop filtration system is not typically needed if the makeup water is of good quality. A side stream filtration system may provide adequate solids removal at a smaller capital cost. The operational aspect of filter monitoring and change-out frequency must be considered and a specific maintenance program established.

- **Scale: a deposition of water-insoluble constituents, formed directly on the metal surface.** These substances change from a soluble state in water to an insoluble state on the metal surface. Scale is typically calcium carbonate. Keeping the hardness value in check will prevent the most common scaling problems. Scale-forming potential is a function of temperature. The chemistry needs to be considered at the warmest temperature in the loop. Scale is generally not a major problem if the loops do not allow evaporation of water. Evaporation or water vapor transfer and subsequent concentration of the chemistry can occur in vented expansion tanks, as well as through fittings and elastomers (gaskets, etc.) in the system. The pH can be an issue here where lower pH values tend to suppress scale formation but do add to the corrosion potential.
- **Microbiological activity: basic organisms such as aerobic bacteria, anaerobic corrosive bacteria, fungi, and algae.** Bacteria cause slime and destruction of nitrite. Bacteria are also an issue with pitting. Specific bacteria can greatly increase the risk of pitting, attacking piping systems at joints and high stress locations. Anaerobic bacteria can live without the presence of oxygen and are generally absent in water with a high pH. Take remedial action when you detect 10 organisms/mL or greater. Aerobic bacteria live and grow in the presence of oxygen, often an indication of slime that can foul equipment. Take remedial action when you detect 1,000 organisms/mL or greater.

Note: Suspended solids and turbidity can be an indication that corrosive products and other contaminants are collecting in the system. Excessive amounts may indicate corrosion, removal of old corrosive products by a chemical treatment program, or the contamination of the loop by another water source. Suspended solids at high velocity can abrade equipment. Settled suspended matter of all types can contribute to pitting corrosion (deposit attack). Similarly, there may be ions present that may also cause these same issues. Some examples are:

- The presence of copper may be an indication of increased copper corrosion and the need for a higher level of copper corrosion inhibitor.
- Excessive iron is an indication that corrosion has increased, existing corrosion products have been released by chemical treatment, piping has been added to the secondary loop, or the iron content has increased in the replacement water.
- The presence of manganese is also important if its concentration is greater than 0.1 ppm.
- Where water-softening equipment is deployed, a total hardness of 10 ppm or greater indicates that the hardness is bypassing the softener, that the softener regeneration is improper, or that some contamination from another system is present, such as a cooling tower or city water.
- The presence of sulfates is often an indication of a process or water tower leak into the TCS loop. High sulfates contribute to increased corrosion because of their high conductivity.

5.1.1.5 Wetted Material Requirements. The CHWS loop of the CDU permits the following material set. The chemicals that are added need to be compatible with all of the loop materials:

Copper Alloys: 122, 220, 230, 314, 360, 377, 521, 706, 836, 952

Polymer/Elastomer:
- Acrylonitrile butadiene rubber (NBR)
- Ethylene propylene diene monomer (EPDM)
- Polyester sealant (anaerobic)
- Polytetrafluroethylene (PTFE)

Solder/Braze
- Solder alloy: 50-50 lead-tin
- Solder flux: Alpha OA-63, Alpha OA-64
- Braze filler Metal: BCuP-3 or BCuP-5
- Braze flux AWS Type 3A

Stainless Steel
- 300 series
- 400 series

Carbon steel

Polypropylene or Polyethylene

5.1.2 Piping Considerations

The placement of the piping for connecting the facility liquid cooling system to the liquid cooling system of the rack is an important consideration in the total design of the data center. This should be considered early in the design phase of a new data center but can be retrofitted into existing data centers. Both are taken into consideration in the layouts proposed in order to minimize cost of initial construction or upgrades to existing data centers. Care must be taken to avoid placement of any cooling piping that may interfere or reduce airflow to the datacom equipment.

Figure 5.4 depicts one option of installing water cooling into a data center and distributing it to liquid-cooled electronic racks. In this case the CDUs are being fed by chilled-water piping that resides on the perimeter of the data center. These CDUs are also located near the perimeter of the data center. Location of the CDUs near the perimeter would concentrate any leaks from the CHWS loop in this area and permit focus of leak detection and containment for this area. By locating the CDUs in this area, the impact on the air distribution that may be required for the racks is minimized. Hose and/or piping distribution to the racks from the CDUs can be laid out below the raised floor and below the racks parallel to the rows (of racks). Also, any

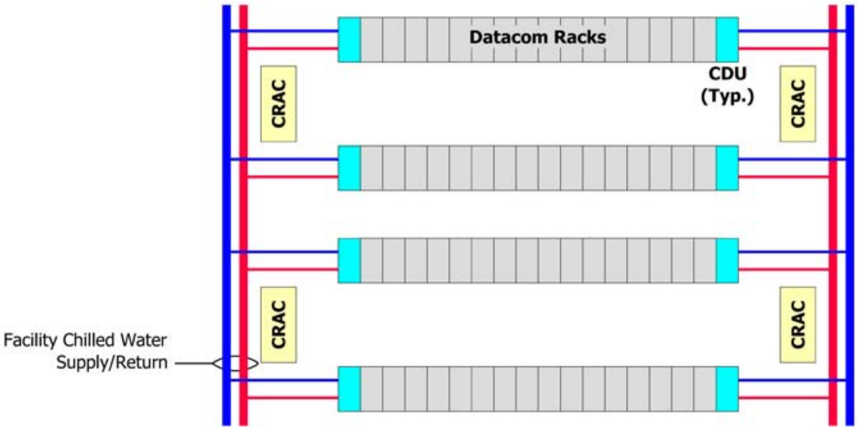

Figure 5.4 Location of CDU units in data center—Option 1.

valves, strainers, or instrumentation in the piping may be easily accessed for operation and maintenance purposes, thereby eliminating risk of accidentally unplugging or harming any communication cables near the racks of equipment.

As an alternative to the CDU units that provide conditioned water at the proper temperature and quality to the electronic racks, these units could be just distribution units that distribute water from the chilled- or process-water lines located at the perimeter of the data center as shown in Figure 5.4. In this case no heat exchange occurs in the distribution unit.

For installation with multiple systems on a closed loop, the installation should include an expansion tank to facilitate the removal of air from the chilled-water lines. If the chilled-water piping does not contain an expansion tank, the chilled-water lines should be vented when bringing a new system online. Pressure on the chilled-water lines must not exceed 100 psi (690 kPa). Chilled-water connections must be accessible.

A slight modification to that shown in Figure 5.4 is shown in Figure 5.5. In this case the CDU units are located against the outer wall of the data center to provide increased control of leak monitoring and detection. In addition, this option may provide improved piping connections between the chilled-water building system and the CDUs. As stated for the option above, the CDUs can be considered to be just liquid distribution units where no heat exchange occurs.

5.1.2.1 Facility Chilled-Water Connections to Datacom Equipment. Datacom equipment racks can be connected to the facility chilled-water systems by either a hard pipe fitting or a quick disconnect attached to OEM flexible hoses. The quick disconnect method is very popular among datacom equipment manufacturers. Each method has its own advantages and disadvantages, which will be discussed further.

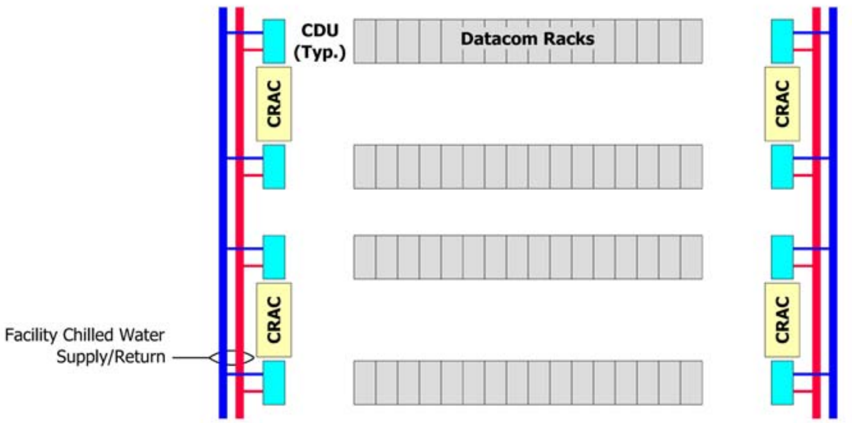

Figure 5.5 Location of CDU units in data center—Option 2.

Using the hard pipe method, connections between the facility chilled-water system and datacom equipment rack can be flanged, screwed, threaded, or soldered, depending on the pipe materials used on both sides of the interface. However, fitting and pipe material must be compatible. The type of pipe material and hard pipe connections will vary by the end user or customer. The end users will have to clarify their requirements or standards to the OEM and design engineer for a successful project.

Most datacom equipment manufacturers requiring connection to the CHWS loops or CDU unit generally provide flexible hoses containing a poppet-style fluid coupling for facilities connection. A poppet-style quick disconnect is an "industrial interchange" fluid coupler conforming to ISO 7241-1 Series B standards. Brass or stainless steel couplings are most commonly used today and must be compatible with the connecting pipe material. If rack loads and flows are excessive, it is recommended that a duplicate set of supply and return lines or hoses be deployed to enhance the fluid delivery capacity of the rack. The design engineer will have to determine if this is necessary during the design. One of the main disadvantages of the quick disconnect is that it has a very large pressure drop or loss associated with it. This pressure loss must be accounted for in all pipe sizing and pump selection procedures. The design engineer must consult with the coupling OEM for exact pressure losses for specific project or system.

Finally, the interface must be properly insulated to prevent any condensation formation. Insulation material should be the same before and after the interface.

The supply and return piping before and after the interface should be properly labeled. This is critical to prevent human error by accidentally switching the two lines.

When using quick disconnects, it is suggested to mix the sockets and plugs between supply and return lines to key them against cross-connection at rack installation.

Other mechanical specialties should be included in the chilled-water piping prior or upstream of the interface. Isolation valves such as ball, gate, or butterfly valves should be installed upstream of the interface to isolate the connection from the building chilled-water system to perform maintenance on the disconnect or to replace it with one of a similar size or larger should the rack load or functionality change. In addition, a strainer should be installed after the isolation valve to catch any particles in the piping system that could damage any of the rack equipment or clog the heat transfer devices. A balancing valve should be provided after the isolation valve to adjust the flow in the building chilled-water system and help push water to other branches within the piping circuit. Finally, any monitoring or controlling instrumentation such as chilled-water pressure, temperature, and flow should be installed at the interface to ensure adequate water conditions exist for proper rack heat exchanger operation.

5.1.3 Electrical Considerations

The electrical power sources and connections serving the CDU units (Figures 5.1 and 5.2) generally must be highly reliable in nature. Liquid-cooled computing applications generally require continuous cooling to ensure continuous operation of the computers. Liquid-cooled cabinets (with built-in heat exchangers) typically require continuous operation of the fans and their control system within the cabinet and the water-circulating pumps that provide the cooling water to the cabinet heat exchanger. The fans, pumps, and their control system are typically powered through a built-in power supply with provision for input from two sources. From this power supply, internal transformers and wiring split the circuits as necessary to support the multiple fans and control system provided with the unit. In the event that one of the power supplies becomes unavailable, internal switching allows all internal electrical components and systems to be fed from the remaining active power supply. Typically, at least one of the feeds to the cabinet will be backed up by an uninterruptible power supply (UPS). The UPS serves to provide continuous power to the cabinet when utility power is unavailable. Liquid-cooled cabinets are likely to be powered from the same UPS system that serves the computers themselves. It is expected that in facilities requiring the highest level availability, the cabinet power supplies will be supported by UPS, and that maintenance bypasses will be provided to ensure that both cabinet supplies are available during maintenance of one of the power paths. The maintenance bypass helps ensure that the internal transfer switch is not relied upon for continued operation during maintenance events. Of course, in all cases, the manufacturer's guidelines for the connections should be followed.

Chilled-water-cooled cabinets generally rely upon chilled water provided by the building infrastructure. Regardless of whether or not the chilling source is from a central building chiller or a dedicated chiller, continuous chilled-water supply will likely be necessary to ensure continued cooling. Electrically, to ensure water supply to the cabinet heat exchanger, the circulating pumps and their control systems should also be powered by a UPS. This could be either a separate UPS dedicated to the mechanical plant or the same UPS that powers the computers, depending upon the electrical systems to be provided.

Modular cooling CDU units (Figure 5.2) are usually provided by the computer manufacturers and matched to the computers that they are expected to cool. As with the liquid-cooled cabinets, these units must be installed and wired in accordance with the manufacturers' instructions. These units typically contain the heat exchanger, pumps, and controls necessary to support the associated computer. Without continuous operation of this equipment, the associated computer cannot be expected to operate. Similar to the computers themselves, the modular cooling units are also fitted with built-in power supplies with provision to be powered from two separate sources. From an electrical standpoint, the power feeds to the modular cooling units should have the same reliability as the power feeds to the computer. It is expected that both the power supplies to the modular cooling unit and to the computer will be backed up by a UPS.

5.1.4 Monitoring

Depending upon the nature and configuration of concept and design of each application, there could be numerous variations of monitoring schemes and parameters to be monitored in a liquid-cooled system. The primary goal is to measure the critical parameters of the liquid systems supporting cooling of the computers (flow, pressure, temperature, etc.) and communicate those values to those responsible for the operation of the equipment. Before considering the specific parameters and associated nuances, let us first consider some high-level communication standards.

Critical facilities typically use sophisticated building automation systems (BAS), which are sometimes also referred to as energy management systems or facility monitoring systems. These systems are undergoing an industry-wide transformation as they migrate away from a strictly proprietary control language to a Web-based communication strategy. Likewise, the initial attempts to develop some open protocol standards by way of BacNet, LonWorks, and manufacturer-specific integrators (gateways) with vendor lookup tables (VLT) are now being challenged to go a step further to be capable of communicating with standard protocols, such as XML, SOAP, etc. This trend and market-driven effort, coupled with the increasing realization that critical infrastructure and IT operations have become "close-coupled" in "real-time," has resulted in a collaborative push to create a mutual standard for communicating data between the traditional silos of IT and facilities services.

A leading organization in this effort is called oBIX, or Open Building Information Xchange (previously called CABA XML/Web Services Guideline Committee), which is a technical committee at the Organization for the Advancement of Structured Information Standards (OASIS). Manufacturers of liquid-cooled hardware, racks, and infrastructure equipment are encouraged to keep abreast of these standards and open protocols and ensure their internal and integral monitoring components and/or systems can communicate easily with both the Enterprise Network (IT) and the critical infrastructure monitoring systems (facilities services), which will both need access to this critical information.

Other recommendations to be considered are that the manufacturers must provide clearly documented specifications that define acceptable ranges, alert values, alarm values, and (if present) any automatic shutdown values and show where, if these values are exceeded, actions are required to protect the hardware and associated applications. Consideration must also be given to the accuracy and maintainability of these monitoring systems and components. For example, thermistors may be more reliable and hold calibration better than thermocouples for the measurement of critical temperatures. Sensors should be housed in drywells if possible to allow removal and replacement without a loss of liquid inventory or an outage.

Where possible, "plug and play" designs should be employed that allow a vendor's product to be easily set up and installed into a pre-existing rack or monitoring system backbone. Computer and server manufacturers should take advantage of the existing network and communications in place for their hardware to communicate with the Enterprise as a means to convey the critical liquid cooling and other environmentally related parameters to the IT departments consistent with other application specific data. They should also include easy means for infrastructure specific systems (BAS) to connect directly to their hardware to retrieve the same environmental data directly.

The incorporation of monitoring standards is being blended, with a good example being the introduction of electrical "smart" power strips. The need to monitor and balance electrical load, especially between redundant power paths, prompted new power strip products capable of not only monitoring the amount of power supplied through each outlet but also having the capability to control each outlet's ON/OFF condition. These intelligent devices can communicate through the existing enterprise LAN/network, thereby eliminating the need to install the separate BAS LAN to each associated cabinet. Many "smart" power strips also offer the ability to add environmental sensors to communicate internal cabinet data with breaker data. By providing a single communication interface between the Enterprise LAN and the BAS LAN; these data can also be linked to the building management system for annunciation of out-of-normal thresholds prior to thermal shutdown conditions. A 24-hour staff is usually necessary to allow adequate response time in these circumstances. If a full-time staff is not present, an added feature could be to allow a remote operator to shut down affected equipment in an orderly process.

5.1.5 Reliability and Availability

Reliability and availability for liquid-cooled racks and electronics are two distinctly different items. By one definition, the *reliability* of a system or subsystem refers to the probability that the system or subsystem will remain operational for a given period of time. Typical metrics for reliability include MTBF (mean time between failure) and L10 life (given as the length of time within which 10% of the devices will fail). *Availability* generally refers to the length of time during which equipment is available for work. One measure of availability is uptime.

This section provides information on the key subsystems that are to be considered when designing for high reliability and availability. Information is also provided regarding a number of events that can impact reliability and availability.

5.1.5.1 Subsystems Affecting Reliability and Availability. Chapter 4 provides a detailed description of the various implementations of liquid-cooled racks and electronics. The reliability and availability of these are highly dependent upon key subsystems and sensors. Following is a list of the key subsystems and sensors to be considered, along with a brief discussion.

- **Pumps**—These cover water, refrigerant, and dielectric pumps for the various implementations. Depending upon the liquid cooling implementation, pumps can be present in all of the coolant loops (CWS, CHWS, TCS, and DECS). Redundancy can be provided in each loop, but this is more difficult and less cost-effective at the individual server level (particularly 1U and single blades). Multiple pumps running in parallel are typically more efficient than one and can provide improved reliability.
- **Fans and blowers**—These are used in many of the implementations of liquid cooling. Examples include (1) servers where only the processors are liquid-cooled and the balance of the servers are air-cooled or (2) enclosed racks where the electronics are air-cooled and the heat is rejected to an air-to-liquid heat exchanger. Fans and blowers are especially critical for enclosed racks, although such systems typically offer $n + 1$ or better redundancy.
- **Controllers**—These can be server-, rack-, or facility-based. The controllers control or interface with rack-based pumps, fans, and power distribution units, as well as the same devices at the facility level. Both hardware and firmware are of concern with regard to the controllers.
- **Leak detection systems**—These may use rope or other types of sensors to detect water (most critical), refrigerant, and dielectric leaks at the server, rack, and facility levels. Such systems can interface with rack- or facility-based controllers and can bring varying levels of the datacom facility down in the event of a leak. They can also send various types of alarms out in the event of a leak. Such systems may also be implemented with water shutoff valves that allow isolation of the racks.
- **Heat exchangers**—These can be found at the server, rack, and facility levels. Fouling is of particular concern, and the type of fouling varies depending upon the type of liquid that is passing through the heat exchanger. Pressure drop through these devices should be minimized to reduce the pumping power required to pump the liquid through them.

- **Connectors and fittings**—These can be found at the server, rack, and facility levels. High-quality fittings are necessary to ensure leak-free joints. Further details are provided in Section 5.1.2.1.
- **Tubing**—Tubing for liquid-cooled systems is found at the server, rack/cabinet, and facility levels. High quality tubing is necessary to ensure leak-free operation.
- **Sensors**—Temperature sensors, pressure transducers, flow sensors, and liquid level sensors are among the key sensors in liquid-cooled systems.

5.1.5.2 Issues Impacting Reliability and Availability. A number of events can lead to reduced system reliability or reduced availability. While there are distinct differences between reliability and availability, the two of these are closely related in some ways. For example, frequent power outages that are accompanied by voltage spikes result in a lack of availability during the outage, as well as equipment (electrical) stress that may eventually lead to equipment failure (reduced reliability). Following is a listing and description of some key items that affect reliability and availability.

This section describes the impact of several water-related events on system reliability and availability, e.g., the disruption of water supply to a rack. For the sake of brevity, only water is discussed. The reader should be aware that similar events that involve other fluids such as refrigerants and dielectrics will similarly impact the reliability and availability of liquid-cooled systems. For example, the disruption of a refrigerant supply to a rack will have an impact similar to the disruption of a water supply to a rack. The reference to water only is not intended to suggest that this is the only or the recommended method of implementation for liquid-cooled systems.

- **Power outages**—For facilities with neither UPS nor backup power generation capability, a power outage will result in an immediate removal of availability of the datacom equipment. This lack of power for the electronics will also result in a lack of power for pumps and blowers at the facility, rack, and server levels. Where UPS or backup power generation is available, temporary (UPS) or reduced power will be available to allow for graceful equipment shutdown or reduced operation.
- **Water supply**—This can encompass partial or complete disruption of a water supply or variations in temperature of the water. Changes in the supply pressure (and flow rate) and temperature of the water will impact the capacity to remove heat from the electronics and, in turn, will impact the availability of the rack.
- **Water hammer**—As liquid-cooled solutions are propagated throughout datacom centers, increasing numbers of rack-dedicated water shutoff valves (see "Leak Detection Systems" in Section 5.1.5.1) will be deployed. Improper operation of any of these can lead to water hammer throughout the facility, which, in turn, can lead to damage at the server, rack, and facility levels. Other sources of water hammer should also be considered, and the system design should account for this.

- **Water flow balancing**—Implementation of hundreds of liquid-cooled racks will lead to a much more complicated (than the current norm) flow network in the datacom facility. Improperly designed rack-level cooling loops can lead to improper water flow distribution to racks.
- **Water quality**—Water quality is impacted by a number of items, including the presence of suspended solids, bacterial contamination, and incorrect water pH level. The primary subsystems impacted are:
    - **Heat exchangers**—Improperly maintained water quality will lead to heat exchanger fouling (see also Section 5.1.1.4 for additional discussion of water quality issues), which, in turn, will lead to reduced cooling capacity and reduced availability. In many cases, heat exchangers can provide a single point of failure, and designs that will minimize the impact of these failures should be considered.
    - **Pipes**—Improperly maintained water quality will lead to fouling of the supply and return piping. Pipe fouling could manifest itself in reduced flow in cross sections and an undesirable increase in pressure drops. In many cases, piping or hose can provide a single point of failure, and designs that will minimize the impact of these failures should be considered.
    - **Connectors and fittings**—Improperly maintained water quality will lead to fouling of connectors and fittings. This may lead to challenges during repair and maintenance or leaks in between servicing.
- **Condensation**—Facility water can be supplied as low as 39°F (4°C), which is below the ASHRAE-recommended room dew-point guideline of 63°F (17°C) for Class 1 Enterprise Datacom Centers (ASHRAE 2004b). Equipment vendors are aware of this and are taking this into account in their designs. If low operating temperatures are expected, careful consideration of condensation should be exercised. It is suggested that a CDU with a heat exchanger be employed to raise the coolant temperature to at least 18°C to eliminate condensation issues.
- **Seismic events**—Seismic events will lead to compromising of joints if there is inadequate stress relief at such joints. As an example, the joints between a facility water loop and a rack could be damaged during a seismic event if flexible piping is not utilized. Facilities using seismic isolators with movable rack foundations should ensure a sufficient service loop of flexible tubing to avoid stressing the cooling connections during seismic events.
- **Operator error**—As with any type of cooling system, operator error should always be considered. With a liquid-cooled system, inadvertent disruption of the water supply to a rack will lead to a shutdown (and possible damage) of the rack. Facility design should accommodate such a scenario. Interlocks or automated controls should be considered along with checklists and procedures.

5.1.5.3 Recommendations. Liquid cooling represents a new paradigm for datacom center operators. As with any new approach to doing things, education will play a large role in the successful implementation of liquid cooling. In general, most of the general rules for current air-cooled implementations apply. Datacom center operators should have cooling contingency plans; implement cooling system redundancy; deploy critical subsystems, such as pumps that have high reliability and availability; and should place subsystems such as pumps (water, refrigerant, dielectric) and rack-based cooling systems on UPS.

5.1.6 Commissioning

Commissioning at the interface is an essential part of the project. Since the interface is the boundary (CHWS shown in Figure 1.1) between the facility's cooling systems and the computer cooling systems, it is imperative that all devices function properly and that controls work adequately, predictably, and consistently under both normal and fault conditions. In addition, it is vitally important that the facility's cooling system provide the proper flows, pressures, and temperatures required by the computer/rack manufacturer. Commissioning starts very early in the design phase of the project. All requirements for the datacom equipment should be provided to the design engineer, who then incorporates these requirements into the design. During commissioning, tests and measurements should be performed on the various components, as well as the integrated system, to ensure they function properly and to ensure the datacom equipment, environment, and supporting utilities meet specifications and respond to anomalies as expected.

Depending on the products and configurations, some typical devices or components found at the interface may include valves, sensors, strainers, and quick disconnect couplings. The following list provides examples of tests and/or procedures that could be included in commissioning liquid-cooled rack deployments:

- Valves provide proper shutoff and closure.
- Actuators fully open and close the valve and are accurately monitored by the BAS system (including that the automated valve is mapped to the proper location on any BAS graphic).
- Strainers can be properly drained and screens removed for cleaning.
- BAS sensors provide accurate readings.
- Flowmeters are installed according to manufacturers' recommendations (e.g., in a straight length of pipe with no other fittings or devices around it).
- "Dripless" quick disconnects facilitate installation and removal of new cooling devices.
- System parameters are within the expected range as predicted by the engineered design (e.g., pressure drops, flow rates, water quality, inlet and return temperatures, etc.).

- Establish a written baseline depicting the original conditions and values at time of startup as historical documentation.

In summary, commissioning is important to ensure that all aspects of the deployment are compliant with the original design intent and specifications. Commissioning should follow a formal written plan and should start during the project's design phase. The building owner should consider hiring a commissioning agent to prepare a detailed list of tests and services to be included in the commissioning effort.

5.2 NON-CHILLED-WATER FACILITY SYSTEMS

Chilled water is not the only choice to provide cooling to a data center. Figures 5.6 and 5.7 show alternative solutions, such as direct expansion (DX) equipment with an air -cooled condenser or glycol and water mixture, which uses a dry cooler to cool the solution and return it to the refrigeration equipment. The benefits and design suggestions will be discussed in this section. Further technical information can be found in Chapter 38 in the *2004 ASHRAE Handbook—HVAC Systems and Equipment*.

5.2.1 Air-Cooled Condensers

Air-cooled condensers should be placed in a area that will ensure an adequate air supply, since outside ambient air is the cooling medium. In addition, they should be located in a clean air area that is free of dirt, hot air, steam, and fume exhausts. Restricted airflow through the condenser will reduce operating efficiency of the unit and can result in high head pressure and loss of cooling. As such, units should be located no closer than 5 ft from a wall, obstruction, or adjacent unit. Always refer to the manufacturers' recommendations for maintenance accessibility and location recommendations. A direct expansion system may be less expensive to install than a chilled-water system, depending on the system size, capacity, and growth potential. In addition, a DX system does not require domestic water storage for reliable operation, since cooling towers are not used to condense the working fluid. DX systems can be used to cool an entire data center or provide localized cooling to "hot spots" within the data center.

5.2.2 Refrigerant Piping

The efficiency and reliability of a direct expansion system can hinge on the piping that interconnects the refrigerant-condensing and air-handling sides of the system. Operational difficulties will be readily apparent if the interconnecting piping is not designed and installed properly. The following guidelines will help eliminate or minimize any operational difficulties with interconnecting piping:

- Design a simple and direct layout that reduces the amount of system refrigerant. Route the suction line from the evaporator to the compressor by the shortest path.

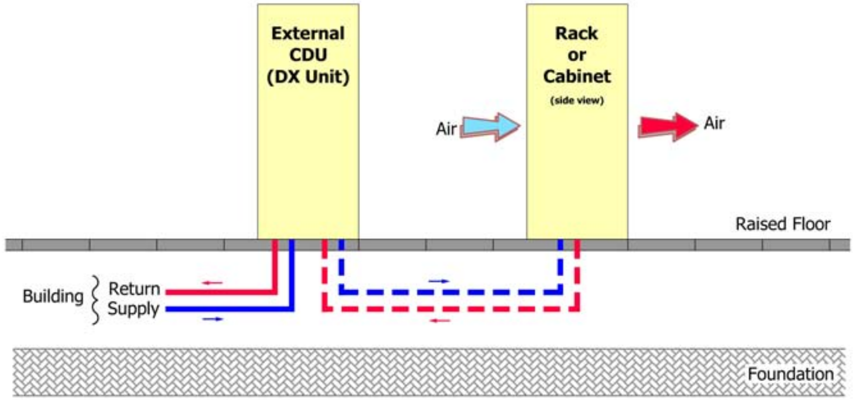

Figure 5.6 CDU (DX unit) supplying coolant to rack or cabinet.

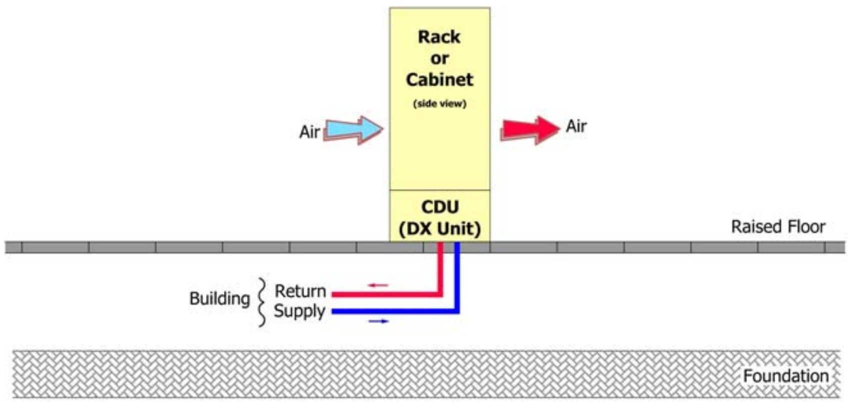

Figure 5.7 Modular CDU (DX unit) within rack or cabinet.

- Limit the overall line length, including the vertical suction and liquid risers. Enough subcooling may be lost as refrigerant travels up the liquid riser to cause flashing.
- Use different pipe sizes for horizontal and vertical lines to make it easier to match line pressure drop and refrigerant velocity to suction-line requirements.
- Properly size the suction line to ensure proper oil entrainment and avoid sound transmission.
- Riser traps are unnecessary. If the riser is properly sized to maintain velocity, adding a trap only increases the suction-line pressure drop.
- Double suction risers may be unnecessary due to the type of compressor selected.
- Eliminate any foreign matter in the system and install suction filters.
- Provide a 1 in. pitch toward the evaporator for every 10 ft of run to prevent any refrigerant that condenses in the suction line from flowing to the compressor when the unit is offline.
- Use insulation on the suction and liquid lines if moisture condensation or dipping causes a problem.
- Select the smallest practical liquid line size for the application. Limiting the refrigerant charge improves compressor reliability.
- The liquid line should include a replaceable filter drier to permit proper system cleanup and removal. The unit should be inspected or changed whenever it is serviced or opened up.
- A moisture-indicting sight glass should be added to permit a visual inspection of the liquid column for bubbles.
- Solenoid valves are required to prevent the liquid refrigerant from filling the evaporator when the compressor stops and will prevent slugging when the compressor restarts. They also prevent siphoning, which could allow an elevated column of liquid to overcome the gas trap and flow back into the compressor.
- Hot gas bypass valves should be considered for scalability. The use of hot gas bypass valves will allow minimum run times and prevent short-cycling when the initial equipment does not develop sufficient load to allow the compressors to run.

6

Liquid Cooling Infrastructure Requirements for Technology Cooling Systems

The interfaces that occur between the rack-level cooling hardware deployed to enhance rack-level thermal management or extend data center thermal management capacity are described here. The position of this interface is shown in Figure 6.1.

6.1 WATER-BASED TECHNOLOGY COOLING SYSTEM

One example of this loop is displayed in Figure 6.1 (same as Figure 4.4); other examples are shown in Figures 4.5–4.14. Since these solutions that remove the heat load near or at the rack are designed by companies that control the details of the liquid and materials within this loop, only some broad guidelines can be given in this section.

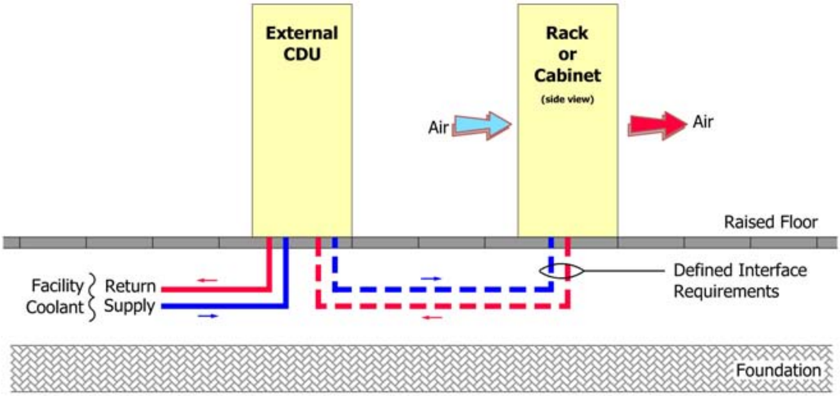

Figure 6.1 Combination air- and liquid-cooled rack or cabinet with external CDU (same as Figure 4.5).

One implementation could have the electronics air-cooled, with the coolant removing a large percentage of the waste heat via a rear door heat exchanger or a heat exchanger located above the rack. Another implementation would include a totally enclosed rack that uses air as the working fluid and an air-to-liquid heat exchanger. Another would be with the coolant passing through cold plates attached to processor modules within the rack. The CDU can be external to the datacom rack, as shown in Figure 6.1, or within the datacom rack, as shown in Figures 4.6 and 4.7.

6.1.1 Operational Requirements

In most cases the liquid that is supplied to the rack or cabinet will be above the room dew point. The ASHRAE book *Thermal Guidelines for Data Processing Environments* (ASHRAE 2004b) specifies a maximum dew point of the room for a Class 1 environment of 63°F (17°C). Two techniques can be used to maintain liquid supply temperatures above the dew point—one is to set the control point of the liquid above the Class 1 environmental specification of 17°C or have a supply temperature controlled such that it is adjusted a set amount above the measured room dew point.

The maximum allowable water pressure for the TCS loop should be 100 psig (690 kPa) or less.

6.1.2 Water Flow Rates

Water flow rates for the TCS loop are set by the manufacturer of such cooling equipment. Figure 6.2 shows the flow rates for given heat loads and given temperature differences. Temperature differences typically fall between 9°F and 18°F (5°C and 10°C).

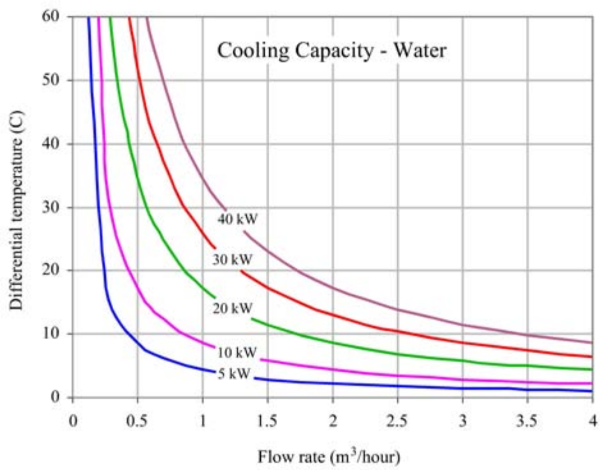

Figure 6.2 Water flow rates for TCS loop and for constant heat load.

6.1.3 Velocity Considerations

The velocity of the water in the rack's liquid cooling loop piping must be controlled to ensure that mechanical integrity is maintained over the life of the system. Velocities that are too high can lead to erosion, vibration, and water hammer. Lower velocities lead to lower pressure drop and lower pumping power required to transport the liquid. Table 6.1 provides guidance on maximum chilled-water piping velocities. Any flexible tubing should be maintained below 5 ft/s (1.5 m/s).

6.1.4 Water Quality/Composition

The quality of cooling water in a technology cooling system (TCS) loop is critical for the performance and the longevity of the datacom equipment. Cooling water of poor quality can cause adverse effects in a water system, such as reduced cooling capacity, increased energy consumption, and premature equipment failure.

Table 6.2 gives the recommended specification for the TCS liquid coolant. If water is outside these ranges, it does not mean the system will have water-quality-driven issues. In fact, water well outside these ranges has been used in cooling loops successfully. The main issue in these other loops is a much greater requirement for, and emphasis on, the water chemistry treatment program. Meeting the proposed water quality specifications will make the required water chemistry program much simpler.

6.1.4.1 Replacement Water Guidelines. Water in the TCS loop should be changed as needed. If the water exceeds the limits of the values in Table 6.2 or exceeds the limits of the established water chemical treatment plan, it should be replaced. Generally a partial replacement will be sufficient to regain proper cleanliness levels.

The TCS loop makeup water source is important in starting with a clean system and maintaining it. A system using the CWS loop as the feed source can be successful; however, as mentioned earlier, the water treatment program's efficacy is critical to the success of that loop. Typically, the TCS loop should be filled with soft, deionized, distilled, or reverse-osmosis product water.

Table 6.1 Maximum Velocity Requirements

Pipe Size	Maximum Velocity (fps)	Maximum Velocity (m/s)
>3 in.	7	2.1
1.5 to 3 in.	6	1.8
<1 in.	5	1.5
All flexible tubing	5	1.5

Table 6.2 Water Quality Specifications—TCS Cooling Loop

Parameter	Recommended Limits
pH	7 to 9
Corrosion inhibitor	Required
Biocide	Required
Sulfides	<1 ppm
Sulfate	<10 ppm
Chloride	<5 ppm
Bacteria	<100 CFU/mL
Total hardness (as $CaCO_3$)	<20 ppm
Conductivity	0.20 to 20 micromho/cm
Total suspended solids	<3 ppm
Residue after evaporation	<50 ppm
Turbidity	<20 NTU (nephelometric)

6.1.4.2. Special Problems in a TCS loop. Because there is no regular blowdown (sediment purge) from a TCS loop, strainers or side-stream filters may be needed to remove debris that exists in the system.

In any water loop, bacteria will develop in “dead legs” (unused runs of piping that have been isolated with valves from the rest of the system) where there is no water flow. Contamination may result when “dead legs” are reconnected to the system; biocides can help reduce the impact but cannot prevent this. Consider the operational use of the liquid system during design. If there are long runs of piping to electronics that may possibly be isolated for long periods of time (a week or more), it may be worth establishing a bypass or some type of minimum flow rate to avoid this potential problem. Regardless of the nature of the water, any biocide or chemical treatment will break down over time in a dead leg, allowing a proliferation of bacteria in the loop.This problem can be exacerbated by underfloor piping systems, where the possibility of “out of site, out of mind” operations can leave unused piping isolated for long periods of time. If a section of pipe is no longer needed, it may be better to remove it or isolate it and drain it rather than leave it charged with water.

The level of bacteria, which may indicate system contamination, is lower in the TCS loop than in the CWS loop. Although the cooling tower can operate effectively with 100,000 or more organisms per mL, you should take corrective action immediately when bacteria counts exceed 1,000 CFU/mL in the TCS loop.

6.1.4.3 Water Treatment. Before any new computer system is placed into operation, you should flush the loops thoroughly to remove as much suspended material and debris as possible. A chemical detergent cleaning is also desirable. It is important to ensure the detergent residue is rinsed away prior to filling the loop for operation.

To avoid loop problems later, you should seek the advice of a water treatment specialist early in the design stage of your system and diligently follow the program that the specialist recommends.

6.1.5 Wetted Material Requirements

This section describes the recommended materials for use in supply lines, connectors, manifolds, pumps, and any other hardware that makes up the TCS loop at your location.

- Copper
- Brass with less than 15% zinc content
- Stainless steel—304 L or 316 L
- Ethylene propylene diene monomer (EPDM) rubber—peroxide cured

Materials to Avoid in the TCS Loop

The following materials must never be used in any part of your water supply system.

- Oxidizing biocides, such as chlorine, bromine, and chlorine dioxide
- Aluminum
- Brass with greater than 15% zinc (unless corrosion inhibitor is added to protect high zinc brass)
- Irons (nonstainless steel)

While aluminum is an excellent heat transfer material, is lightweight, and is low-cost, it can be problematic in a closed-loop system containing copper. While it is possible to inhibit corrosion of both metals in the same system, it is very difficult to do and requires more expensive treatments and a higher level of care and monitoring. It may be simpler, with the pervasiveness of copper in these systems, to avoid the problems by excluding aluminum from the design.

6.1.6 Monitoring

Consideration should be given to providing for appropriate monitoring of critical parameters consistent with typical best practices for any mission critical system. For additional information and guidance, refer to Section 5.1.4, “Monitoring.”

6.2 NON-WATER-BASED TECHNOLOGY COOLING SYSTEM

The most prevalent liquids used in the TCS loop other than water are refrigerants and dielectrics. The refrigerants are generally pumped as single-phase liquid to the datacom racks. Once the liquid reaches the rack, a phase change may occur to achieve high heat transfer. Of course, with these systems the CDU is a specially designed unit that provides the heat exchange and pumping components that are specifically designed for the specific refrigerant or dielectric. In addition, the piping and/or hose components that are used to transport the fluid to the rack are of a special design to eliminate any leak potential and prevent failures.

6.2.1 Operational Requirements

In most cases the liquid that is supplied to the rack or cabinet will be above the room dew point. The ASHRAE book *Thermal Guidelines for Data Processing Environments* (ASHRAE 2004b) specifies a maximum dew point of the room for a Class 1 environment of 63°F (17°C). Two techniques can be used to maintain liquid supply temperatures above the dew point—one is to set the control point of the liquid above the Class 1 environmental specification of 63°F (17°C) or have a supply temperature controlled such that it is adjusted a set amount above the measured room dew point.

6.2.2 Liquid Requirements

The flow rates, velocities, and quality levels of the refrigerant or dielectric will be specific to the design provided by the manufacturer of the cooling system.

6.2.3 Wetted Material Requirements

The materials required for the liquid loop will be specified by the manufacturer of the cooling system.

References and Bibliography

ASHRAE. 2004a. *2004 ASHRAE Handbook—HVAC Systems and Equipment*, Chapters 13, 35, 36, 38, and 39. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE. 2004b. *Thermal Guidelines for Data Processing Environments*. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE. 2005a. *Design Considerations for Datacom Equipment Centers.* Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE. 2005b. *Datacom Equipment Power Trends and Cooling Applications*. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

Baer, D. 2004. Managing data center heat density. *HPAC Engineering* 76(2):44–47.

Beaty, D. 2004. Liquid cooling: Friend or foe. *ASHRAE Transactions* 110(2):643–52.

Beaty, D., and R. Schmidt. 2004. Back to the future: Liquid cooling data center considerations. *ASHRAE Journal* 46(12):42–46.

Belady, C., and D. Beaty. 2005. Data centers: Roadmap for datacom cooling. *ASHRAE Journal* 47(12):52–55.

Belady, C. 2001. Cooling and power considerations for semiconductors into the next century. *Proceedings of the International Symposium on Low Power Electronics and Design, IEEE, Huntington Beach, CA,* pp. 100–105.

Chu, R.C. 2004. The challenges of electronic cooling: Past, current and future. *Journal of Electronic Packaging*, *Transactions of the ASME* 126(4):491–500.

Delia, D.J., T.C. Gilgert, N.H. Graham, U. Hwang, P.W. Ing, J.C. Kan, R.G. Kemink, G.C. Maling, R.F. Martin, K.P. Moran, J.R. Reyes, R.R. Schmidt, and R.A. Steinbrecher. 1992. System cooling design for the water cooled IBM Enterprise System/9000 processors. *IBM Journal of Research and Development* 36(4):791–803.

Kakac, S., and H. Liu. 2002. *Heat Exchangers: Selection, Rating and Thermal Design*. CRC Press, Inc.

Kurkjian, C., and J. Glass. 2005. Air-conditioning design for data centers—Accommodating current loads and planning for the future. *ASHRAE Transactions* 111(2):715–24.

Pautsch, G. 2001. An overview on the system packaging of the CRAY SV2 supercomputer. *Proceedings of IPACK'01*, pp. 617–24.

Schmidt, R. 2005. Liquid cooling is back. *ElectronicsCooling* 11(3):34–38.

Schmidt, R., and B. Notohardjono. High end server low temperature cooling. *IBM Journal of Research and Development* 46(6):739–51.

Schmidt, R., R. Chu, M. Ellsworth, M. Iyengar, D. Porter, V. Kamath, and B. Lehman. 2005. Maintaining datacom rack inlet air temperatures with water cooled heat exchangers. *Proceedings of the Pacific Rim/ASME International Electronic Packaging Technical Conference (INTERpack), San Francisco, CA, July 17–22, Paper IPACK2005-73468.*

Shah, R.K., and D.P. Sekula. 2003. *Fundamentals of Heat Transfer Design*. New Jersey: John Wiley & Sons, Inc.

Stahl, L., and C. Belady. 2001. Designing an alternative to conventional room cooling. *Proceedings of the International Telecommunications and Energy Conference (INTELEC), Edinburgh, Scotland*, October.

Trane. 2001. As equipment evolves so must piping practices... Split systems and interconnecting refrigerant lines. *Engineers Newsletter* 27(4).

Glossary of Terms

aerobic bacteria: bacteria that live and grow in the presence of oxygen, often an indication of slime that can foul equipment; take remedial action when you detect 1,000 organisms/mL or greater.

a***ir cooling:*** the case where only air must be supplied to an entity for operation.

*a**ir-cooled rack:*** the case where only air must be provided to the rack or cabinet for operation.

air-cooled datacom equipment: the case where only air is provided to the datacom equipment for operation.

air-cooled electronics: the cases where air is provided directly to the electronics within the datacom equipment for cooling with no other form of heat transfer; if the datacom equipment contains both liquid-cooled and air-cooled electronics, the equipment itself is considered liquid-cooled.

aluminum: a lightweight, silver-white, metallic element.

anaerobic bacteria: bacteria that can live without the presence of oxygen; generally absent in water with a high pH; take remedial action when you detect 10 organisms/mL or greater.

availability: a percentage value representing the degree to which a system or component is operational and accessible when required for use**.**

BAS: building automation system.

BMS: building management system.

CFD: computational fluid dynamics.

CFU: colony-forming unit.

chilled-water system (CHWS): primarily consists of a system between the data center chiller(s) and the CDU; the chilled-water system includes the chiller plant, pumps, hydronic accessories, and necessary distribution piping at the facility level; this system typically is at the facility level and may include a dedicated system for the information technology space(s) (see Figure 1.1).

chloride: an indication of water softener regeneration problems if the system chloride level is much higher than the chloride level of the replacement water; increased levels of chloride can increase corrosion and indicate the need for the addition of higher levels of corrosion inhibitors.

CMMS: computerized maintenance management system.

condenser water system (CWS): consists of the liquid loop between the cooling tower and the data center chiller(s); it also is typically at the facility level and may or may not include a dedicated system for the information technology space(s) (see Figure 1.1).

conductivity: a measure of the mineral content in the water; in a nitrite program, high conductivity is generally an indicator of bacterial degradation of the nitrite.

copper: a ductile, malleable, reddish-brown, corrosion-resistant diamagnetic metallic element.

coolant distribution unit (CDU): conditions the technology cooling system (TCS) or datacom equipment cooling system (DECS) coolant in a variety of manners and circulates it through the TCS or DECS loop to the rack, cabinet, or datacom equipment.

corrosion: deterioration of intrinsic properties in a material due to reactions with its environment.

CRAC: computer room air conditioner.

data center: a building or portion of a building whose primary function is to house a computer room and its support areas; data centers typically contain high-end servers and storage products with mission critical functions.

datacom: a term that is used as an abbreviation for the data and communications industry.

datacom equipment cooling system (DECS): an isolated loop within the rack that is intended to perform heat transfer from the heat-producing components (CPU, memory, power supplies, etc.) to a fluid-cooled heat exchanger also contained within the IT rack; this system is limited to the information technology (IT) rack (see Figure 1.1).

dead lets: unused runs of piping that have been isolated with valves from the rest of the system.

dielectric fluid: a fluid that is a poor conductor of electricity.

DX: direct expansion.

economizer, air: a ducting arrangement and automatic control system that allows the cooling supply fan system to supply outdoor (outside) air to reduce or eliminate the need for mechanical refrigeration during mild or cold weather.

economizer, water: a system by which the supply air of a cooling system is cooled directly or indirectly or both by evaporation of water or by other appropriate fluid (in order to reduce or eliminate the need for mechanical refrigeration).

EMCS: energy management and control system.

firmware: data stored in a computer's read-only memory (ROM) or elsewhere in a computer's circuitry that provides instruction for the computer or hardware devices; unlike normal software, firmware cannot be changed or deleted by an end-user and remains on the computer regardless if it is on or off.

fouling factors: added resistance to the transfer heat from the liquid caused by deposits fouling the heat transfer surface.

heat pipe: also defined as a type of heat exchanger, a tubular-closed chamber containing a fluid in which heating one end of the pipe causes the liquid to vaporize and transfer to the other end where it condenses and dissipates its heat; the liquid that forms flows back toward the hot end by gravity or by means of a capillary wick.

hydronic: a term pertaining to water used for heating or cooling systems.

IT: information technology.

iron: a heavy, ductile, magnetic metallic element; is silver-white in pure form but readily rusts.

liquid cooling: the case where liquid must be circulated to and from the entity for operation.

liquid-cooled rack: the case where liquid must be circulated to and from the rack or cabinet for operation.

liquid-cooled datacom equipment: the case where liquid must be circulated to and from the datacom equipment for operation.

liquid-cooled: defines the cases where liquid must be circulated directly to and from the electronics within the datacom equipment for cooling with no other form of heat transfer.

manganese: a gray-white or silvery, brittle metallic element, which resembles iron but is not magnetic; important only if manganese is present in concentrations greater than 0.1 ppm in the replacement water.

mg per liter: equivalent to ppm.

molybdate: a commonly used corrosion inhibitor.

MTBF: mean time between failure.

nitrite: a commonly used corrosion inhibitor.

NTU: nephelometric turbidity unit.

OASIS: Organization for the Advancement of Structured Information Standards.

oBIX: open building information exchange.

pH: a measure of hydrogen concentration, used to determine whether the water has either corrosive or scaling tendencies; pH is a logarithmic scale of concentration of hydrogen ions (H+) as compared to that of pure distilled water that has an equivalent pH of 7.

PPM: parts per million.

rack: frame for housing electronic equipment.

redundancy: often expressed compared to the baseline of *N*, where *N* represents the number of pieces to satisfy the normal conditions; some examples are "*N* + 1," "*N* + 2," "2*N*," and "2(*N* + 1)"; a critical decision is whether *N* should represent normal conditions or whether *N* includes full capacity during offline routine maintenance. Facility redundancy can apply to an entire site (backup site), systems, or components; IT redundancy can apply to hardware and software.

refrigerants: in a refrigerating system, the medium of heat transfer that picks up heat by evaporating at a low temperature and pressure and gives up heat on condensing at a higher temperature and pressure.

reliability: a percentage value representing the probability that a piece of equipment or system will be operable throughout its mission duration; values of 99.9% ("three nines") and higher are common in data and communications equipment areas; for individual components, the reliability is often determined through testing; for assemblies and systems, reliability is often the result of a mathematical evaluation based on the reliability or individual components and any redundancy or diversity that may be employed.

SCADA: system control and data acquistion.

scale: a deposition of water-insoluble constituents, formed directly on the metal surface.

sulfate: an inorganic ion that is widely distributed in nature.

suspended solids and turbidity: a cloudy condition in water due to suspended silt or organic matter.

technology cooling system (TCS): serves as a dedicated loop intended to perform heat transfer from a datacom equipment cooling system into the chilled-water loop; this system would not typically extend beyond the boundaries of the information technology space—the exception is a configuration in which the facility conditioning and circulating unit is located outside the data center (see Figure 1.1).

temperature, dew point: the temperature at which water vapor has reached the saturation point (100% relative humidity).

total hardness: the sum of the calcium and magnesium ions in water.

VLT: vendor lookup table.

Appendix

Survey of Customer Water Quality of Chilled-Water System Loop

Parameter	Recommended Limits	Typical Values Prior to Treatment
pH	7 to 9	5.2 to 9.9
Corrosion inhibitor	Required	–
Sulfides	<10 ppm	7 to 290
Sulfate	<100 ppm	62 to 2780
Chloride	<50 ppm	29 to 3070
Bacteria	<1000 CFU/mL	450 to 7700
Total hardness (as $CaCO_3$)	<200 ppm	80 to 2750
Residue after evaporation	<500 ppm	140 to 13986
Turbidity	<20 NTU (nephelometric)	3.3 to 640

The table above depicts the water quality of a number of computer facilities that were surveyed prior to the installation of systems that required liquid cooling. The facilities that are shown in this table were all able to meet the recommended limits of this document after they implemented water treatment plans.

Index

A

B

C

D

E

F

G

H

L

M

O

P

R

S

T

U

V

W

Data center IT equipment today is predominantly air cooled. However, with rack heat loads steadily climbing, the ability for many data centers to deliver either adequate airflow rates or sufficient chilled air is now being stretched to the limit. These trends in the heat load generated from IT equipment can have detrimental side effects, such as decreased equipment availability, wasted floor space, and inefficient cooling system operation. This situation is creating a need for implementing liquid cooling solutions. The overall goals of the liquid implementations include aspects such as transferring as much waste heat to the facility liquid cooling loop as possible, reducing the overall volume of airflow needed by the racks, and reducing processor temperatures such that increased compute performance can be achieved.

This book on liquid cooling is divided into six chapters and includes definitions for liquid and air cooling as it applies to the IT equipment, describing the various liquid loops that can exist in a building that houses a data center. It also provides the reader an overview of the chilled water and condenser water systems and an overview of datacom equipment cooling options. The book also bridges the liquid cooling systems by providing guidelines on the interface requirements between the chilled water system and the technology cooling system and outlines the requirements of those liquid-cooled systems that attach to a datacom electronics rack and are implemented to aid in data center thermal management.

This book is the fourth in a series of datacom books published by ASHRAE and authored by TC 9.9, Mission Critical Facilities, Technology Spaces, and Electronic Equipment. The other books, listed in order of publication, are *Thermal Guidelines for Data Processing Environments*, *Datacom Equipment Power Trends and Cooling Applications*, and *Design Considerations for Datacom Equipment Centers*.

ISBN: 1-933742-05-4
9781933742052
90430